OPTIMIZATION OF PALLADIUM-CATALYZED *IN SITU* DESTRUCTION OF TRICHLOROETHYLENE-CONTAMINATED GROUNDWATER USING A GENETIC ALGORITHM

THESIS

Manuel Fernández, Captain, USAF

AFIT/GEE/ENV/01M-02

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U. S. Government.

AFIT/GEE/ENV/01M-02

OPTIMIZATION OF PALLADIUM-CATALYZED *IN SITU* DESTRUCTION OF TRICHLOROETHYLENE-CONTAMINATED GROUNDWATER USING A GENETIC ALGORITHM

THESIS

Presented to the Faculty

Department of Systems and Engineering Management

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science in Engineering and Environmental Management

Manuel Fernández, B.S.

Captain, USAF

March 2001

AFIT/GEE/ENV/01M-02

OPTIMIZATION OF PALLADIUM-CATALYZED *IN SITU* DESTRUCTION OF TRICHLOROETHYLENE-CONTAMINATED GROUNDWATER USING A GENETIC ALGORITHM

Manuel Fernández, B.S.
Captain, USAF

Approved:

Dr. Mark N. Goltz (Chairman)	2 Mar 01 date
Lt. Col. Raymond R. Hill (Member)	2 Mar 01 date
Dr. Junqi Huang (Member)	3/2/2001 date

ACKNOWLEDGEMENTS

I really appreciate the efforts of all those who made the completion of this endeavor possible. First, my thesis advisor, Dr. Mark Goltz, who gave me a great deal of technical knowledge in his classes and who provided continued support and feedback through the entire thesis process. Also, Dr. Junqi Huang, who devoted a great deal of time to this thesis. He updated the fate and transport code numerous times to meet my needs and taught me a great deal about the workings of the model and the code. Also, Lt Col Hill, who taught me about GAs and helped me to implement them in this effort. Also, Capt John Christ and 1Lt Craig Garrett, who gave me advice on how to approach the thesis effort and great ideas on how to implement the GA. And, the computer administrators, especially Ron Adams and Jim Gray, both of whom gave me all the computer support I asked for.

Most importantly, I would like to thank my wonderful wife. Her continued support and encouragement over the past eighteen months has been invaluable to my success here at AFIT.

Manuel Fernández

TABLE OF CONTENTS

Page

LIST OF FIGURES

Figure Page

Figure Page

LIST OF TABLES

Table Page

ABSTRACT

Conventional technologies for the treatment of groundwater contaminated with chlorinated solvents have limitations that have motivated development of innovative technologies. One such technology currently under development involves using palladium-on-alumina (Pd/Al) as a catalyst to promote dechlorination. Pd/Al catalyst may be used in-well as part of a re-circulating horizontal flow treatment well (HFTW) system. An HFTW system involves two or more dual-screened wells, with in-well reactors, to capture and treat contaminated groundwater without the need to pump the water to the surface.

In this study, objective and fitness functions, based on system costs and TCE concentration requirements, were developed to optimize a dual-well HFTW system with in-well Pd/Al reactors in a two-aquifer remediation scenario. A genetic algorithm (GA) was coupled with a three dimensional numerical model of contaminant fate and transport to determine optimized HFTW control parameters (well location, pumping rate, and reactor size). The GA obtained a solution within the specified constraints, but the solution was an artificial solution, as contaminated groundwater in one of the two aquifers received no treatment. Based on these results, new objective and fitness functions were developed in an effort to determine the most cost effective solution to remove contaminant mass from the aquifer. The solution arrived at using this approach, while resulting in minimized values of cost per contaminant mass destroyed, produced unacceptably high downgradient contaminant concentration levels. We conclude that by specifying that only two wells could be used in the HFTW system, we overconstrained the problem and that a multi-well HFTW solution is required.

OPTIMIZATION OF PALLADIUM-CATALYZED *IN SITU* DESTRUCTION OF TRICHLOROETHYLENE-CONTAMINATED GROUNDWATER USING A GENETIC ALGORITHM

1.0 INTRODUCTION

1.1 MOTIVATION

Releases of toxic chemicals by industry over the years have greatly burdened our environment and threatened human health and ecosystems. This problem became evident in the 1970s when numerous catastrophes such as Love Canal in New York and The Valley of Drums in Kentucky were discovered. In response, Congress passed the Comprehensive Environmental Response, Compensation, and Liability Act (CERCLA) in 1980 as a means to enforce cleanup of hazardous waste sites (LaGrega et al., 1994). The law authorized the U.S. Environmental Protection Agency (EPA) to coerce responsible parties to clean up contaminated sites. It also allowed the EPA to clean up sites and then receive reimbursement from the responsible parties (US General Accounting Office, 1999). The law is also known as Superfund, a term used to describe the funds designated for cleanup of some abandoned sites. At the time the law was passed, Congress and the American public did not realize the magnitude of the contamination and expected the sites to be cleaned up quickly (LaGrega et al., 1994).

CERCLA established the National Priorities List (NPL), a list which contains the most heavily polluted sites across the nation and which is updated annually by the EPA. Once a facility is assigned to the NPL it must be cleaned up in accordance with CERCLA's procedures and standards. As of December 2000 the NPL contained 1,229 sites (US EPA, 2000), and a survey of the EPA database by the United States General Accounting Office (US GAO) in 1998

showed that another 1,789 sites are potential candidates for the NPL (US GAO, 1998). However, these sites do not represent all the hazardous waste sites because many contaminated sites are not eligible for the NPL. According to the National Research Council (NRC), there are between 300,000 and 400,000 contaminated sites (NRC, 1994).

Of the sites on the NPL, 120 are the responsibility of the Department of Defense (DoD), of which the Air Force is responsible for 37 (DERP, 1998). In DoD, the cleanup of these sites, along with other contaminated sites, falls under the Defense Environmental Restoration Program (DERP). As of 1998, the Air Force was responsible for over 6,038 sites monitored under DERP. The cost for this cleanup is very high; DoD projects expenditures of approximately $1.3 billion yearly through FY01, of which approximately $300 million will be spent by the Air Force.

According to the NRC, the EPA estimates that approximately 80% of the contaminated sites have contaminated groundwater (NRC, 1994). Because groundwater provides around one-third of our drinking water (Masters, 1997), its contamination poses great human health risks, which has lead to much research in groundwater remediation technologies. Trichloroethylene (TCE), a chemical used by industries as a degreaser and solvent, is the most frequently found contaminant at hazardous waste sites (NRC, 1994). TCE has also been used widely at DoD facilities, where numerous industrial processes take place, and groundwater contamination by TCE is also prevalent. TCE is a suspected human carcinogen (Masters, 1997) that may undergo a series of reactions, called reductive dechlorination, which transforms it, successively, into dichloroethylene (DCE), vinyl chloride, ethene, and ethane. Vinyl chloride is known to be a human carcinogen (Masters, 1997), but ethene and ethane are virtually harmless (Maymó-Gatell et al., 1999).

Commonly employed methods used to contain contaminated groundwater plumes are pump-and-treat systems, funnel-and-gate systems, and natural attenuation. Of these, the most common method is pump-and-treat (NRC, 1994), which consists of pumping the contaminated groundwater to the surface, where it is treated using any of a number of different methods such as air stripping, carbon adsorption, or biological treatment (LaGrega et al., 1994). Such systems may incur great costs for pumping and for the disposal of hazardous waste that may be created from the water treatment. Another disadvantage of pump-and-treat systems is that pumping contaminated water to the surface also increases health risks for those who might be exposed to it.

Funnel-and-gate systems are an emerging technology that use zero-valent metals (typically zero-valent iron) in a subsurface trench, the "gate," to remediate water contaminated with chlorinated solvents. The technology is implemented by directing the plume of contaminated groundwater towards the reactive trench using sheet pile walls or other impermeable barriers, the "funnel". However, the technology is limited to certain hydrogeologic conditions. For instance, due to limitations on the depth to which the trench and barriers may be placed, the contamination must be relatively shallow. Also, since the technology is passive, variations in the direction of groundwater flow over time may allow the contaminant plume to bypass the trench. Advantages of funnel-and-gate systems include low maintenance requirements and the fact that the contaminant is destroyed *in situ*, with the attendant cost and safety benefits (Ferland, 2000).

Natural attenuation is the third commonly employed method of containing contaminated groundwater plumes. The US EPA defines natural attenuation as follows (US EPA, 1999):

> The term "monitored natural attenuation" refers to the reliance on natural attenuation processes (within the context of a carefully controlled and monitored clean-up approach)

> to achieve site-specific remedial objectives within a time frame that is reasonable compared to other methods. The "natural attenuation processes" that are at work in such a remediation approach include a variety of physical, chemical, or biological processes that, under favorable conditions, act without human intervention to reduce the mass, toxicity, mobility, volume, or concentration of contaminants in soil and groundwater. These *in situ* processes include biodegradation, dispersion, dilution, sorption, volatilization, and chemical or biological stabilization, transformation, or destruction of contaminants.

Monitored natural attenuation is currently an accepted method to control plumes of fuel hydrocarbons such as benzene, toluene, ethylbenzene, and xylene, which have all been shown to degrade readily in a subsurface environment (Wiedemeier et al., 1998). Unfortunately, chlorinated compounds undergo reductive dechlorination only under specific aquifer conditions and it is the EPA's view that "significant biodegradation of chlorinated solvents sufficient to achieve remediation objectives within a reasonable timeframe are anticipated to occur only in limited circumstances" (US EPA, 1999). The main cost associated with the technology is the monitoring of the site. Additional disadvantages of natural attenuation include longer time period to achieve goals, production of TCE daughter products, and hesitance by regulators and the public to accept the technology (Feng, 2000).

Because of the limitations of these conventional technologies, as well as the desire to reduce cost and risks, new technologies are being developed. One new technology that has the potential to safely and cost-effectively treat chlorinated ethene-contaminated groundwater uses a horizontal flow treatment well (HFTW) system with in-well palladium reactors (Figure 1.1). This system consists of multiple wells, some of which pump in an upflow mode and others in a downflow mode. These wells contain palladium reactors, which when combined with dissolved hydrogen gas, degrade chlorinated contaminants (Ferland, 2000). The upflow/downflow configuration allows for multiple passes of contaminated water through the palladium reactors,

which permits the attainment of very low contaminant concentration downgradient of the treatment system. McCarty et al. (1998) completed a field evaluation using an HFTW system at Edwards AFB, CA. In this study the HFTW system was implemented using bioremediation to treat TCE. McNab et al. (2000) conducted a field experiment at Lawrence Livermore National Laboratories (LLNL) using an in-well palladium reactor in a single dual-screened well to degrade TCE. Both experiments were successful in treating TCE contaminated groundwater. Advantages of an HFTW system using in-well palladium reactors include: (1) reduced risk and cost because the contaminants are not pumped to the surface, (2) decreased concern over formation of harmful products because the chlorinated contaminants are completely dechlorinated, and (3) active control of the contaminant plume (Ferland, 2000). The primary goal of employing *in situ* containment using an HFTW system would be to ensure that contaminant levels downgradient of the treatment system meet regulatory standards. A model to simulate groundwater flow and contaminant transport by an HFTW system has been developed (Christ et al., 1999).

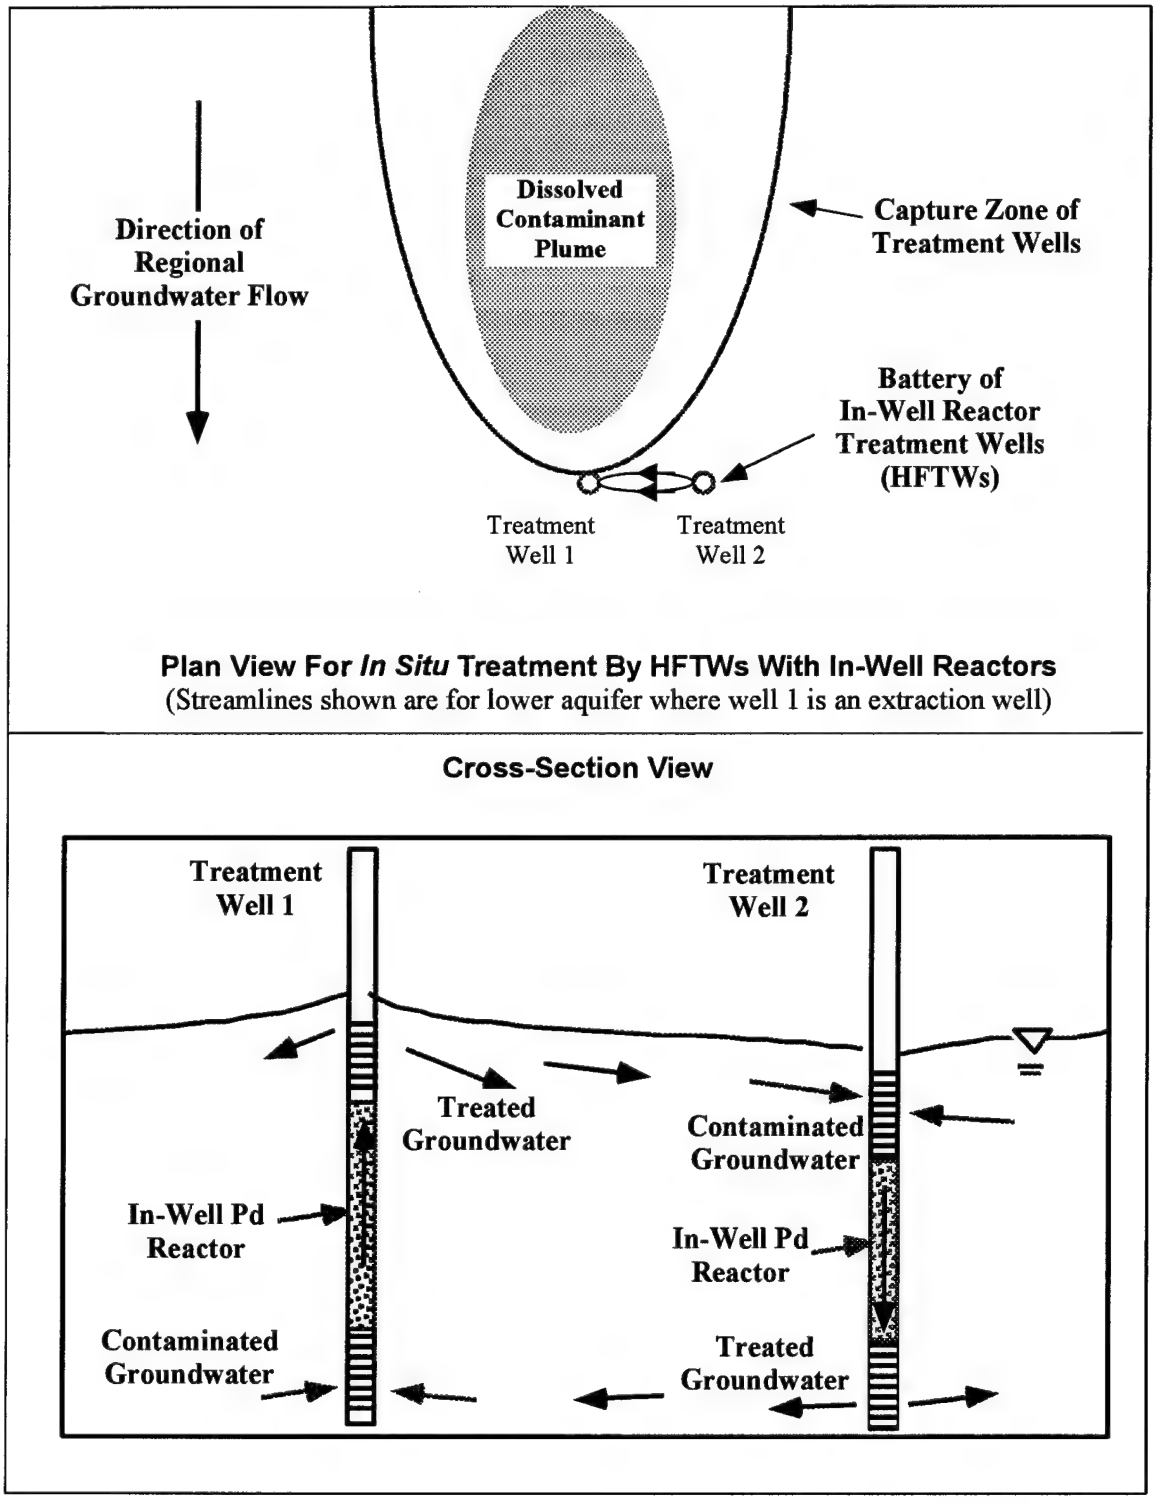

Figure 1.1: Schematic of horizontal flow treatment well (HFTW) system with in-well reactors (Ferland, 2000).

Before such an innovative treatment technology can be widely implemented, project managers must be able to model and predict system performance under specific site conditions. Models are tools that can be used to help designers and project managers understand how a technology will perform under different site and operating conditions as well as quantify the costs of implementation. However, a number of the parameters associated with the technology, such as reactor size, distance between wells, and pumping rate, must be engineered to provide the most effective system. Optimization methods may be used to efficiently determine those parameter values that yield the "best" design. The optimization is performed through the use of an objective function, which is a mathematical representation of the design's objective. The objective function defines the "best" design. In many cases, the objective function will be cost, and the goal of the design will be to minimize cost by varying engineered parameters. Typically, there will also be constraints on the design. That is, for example, the design must minimize cost while assuring capture of the entire contaminant plume and/or attaining a specified contaminant concentration downgradient of the designed system. An optimization algorithm attempts to determine the combination of engineered parameters that minimizes/maximizes some objective function within specified constraints.

Often, the optimization algorithm is used in conjunction with a fate-and-transport model. A fate-and-transport model is a mathematical model that represents the movement of the groundwater and the contaminant in the aquifer. While searching for the optimal solution, the optimization algorithm will specify potential solutions (or sets of parameter values). The fate-and-transport model will be used to determine contaminant distribution in time and space for each of these potential solutions. These distributions are then used by the optimization algorithm to quantify how well the solution meets the objective and/or constraints.

1.2 RESEARCH OBJECTIVES

(1) Develop a methodology for determining how to effectively implement palladium-catalyzed *in situ* remediation under given site conditions.

1.3 SCOPE AND OVERVIEW

This thesis research is limited to the numerical optimization of HFTW systems with in-well palladium reactors to implement *in situ* destruction of TCE-contaminated groundwater. TCE is the contaminant of interest because it is the most frequently found groundwater contaminant at Superfund sites and it is also prevalent at many Air Force bases. The HFTW system with palladium-catalyzed reactors is the technology of interest because of its great potential, as demonstrated in recent field studies, as well as the other advantages of the technology discussed previously.

In the following chapter we will further discuss the problem of TCE-contaminated groundwater, the technology, the fate-and-transport model used to simulate remediation by HFTW systems with in-well palladium reactors, and optimization methods.

2.0 LITERATURE REVIEW

2.1 OVERVIEW

In this chapter, literature pertinent to the optimization of a horizontal flow treatment well (HFTW) system with in-well palladium reactors for degradation of TCE-contaminated groundwater will be reviewed. In the first section, the problem of TCE-contaminated groundwater is briefly described. In the second and third sections, the groundwater treatment technology and models that may be applied to simulate the technology, respectively, are presented. In the fourth section, optimization is defined and general information on the subject is provided. In the fifth section, optimization techniques, including an in-depth look at genetic algorithms, are discussed. Finally, in the last section, applications of groundwater remediation technology optimization are reviewed.

2.2 PROBLEM

TCE is a suspected human carcinogen (Masters, 1997: 183) and it is the most commonly found contaminant in groundwater (NRC, 1994: 26). TCE may undergo a series of reductive dechlorination reactions that transforms it, successively, into dichloroethylene (DCE), vinyl chloride, ethene, and ethane as shown in Figure 2.1. Vinyl chloride is known to be a human carcinogen (Masters, 1997: 183), but ethene and ethane are harmless products (Maymó-Gatell et al, 1999). The EPA has set a drinking water maximum concentration limit, or maximum contaminant level (MCL), of 5 parts per

billion (ppb) or μg/L for TCE and 2 ppb for vinyl chloride (Masters, 1997: 271) because of their threat to human health.

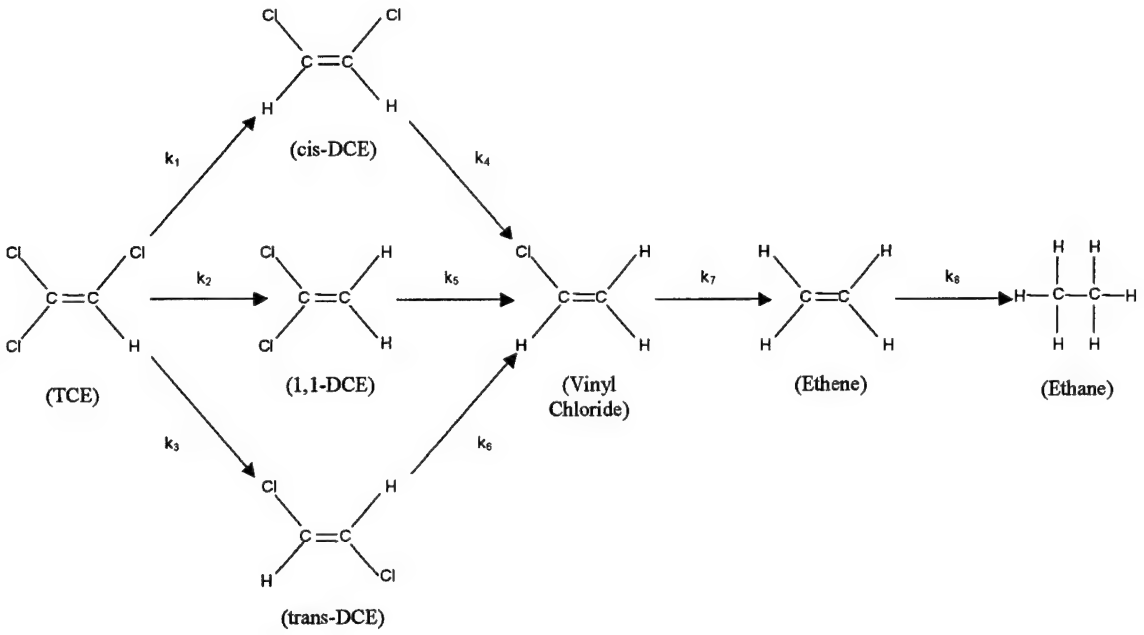

Figure 2.1 Reductive dechlorination of TCE to ethane.

Currently three methods are primarily used to clean up groundwater contaminated with these chlorinated solvents. They are pump-and-treat, natural attenuation, and permeable reactive barriers. However, as discussed in Chapter 1, all these technologies have significant drawbacks that limit their use. Two of these limitations that must be overcome are incomplete dechlorination, which may result in accumulation of vinyl chloride, and extraction of contaminated groundwater to the surface, which increases risk and cost. Innovative technologies that reduce or eliminate these limitations are currently being investigated. One emerging technology that is the subject of recent research uses in-well palladium reactors to catalyze the reductive dehalogenation process.

2.3 TECHNOLOGY

In this section information on the technology is presented. First, palladium-catalyzed reductive dehalogenation is presented through a discussion of laboratory and field experiments. This is followed by a description of a horizontal flow treatment well (HFTW) system.

2.3.1 PALLADIUM CATALYZED REDUCTIVE DEHALOGENATION

A catalyst is a substance that is used in a chemical reaction to change the speed of the reaction. The catalyst itself does not undergo a permanent change during the reaction and it is not consumed (Brown et al., 1997: 521). Palladium-catalyzed reductive dehalogenation is a chemical process where a chlorinated hydrocarbon contaminant (e.g. TCE) is degraded by reacting with dissolved hydrogen gas in the presence of a palladium catalyst. Lowry and Reinhard (1999) performed laboratory experiments that showed what appears to be a direct pathway from TCE to ethane (Figure 2.2) without the formation of any intermediate chlorinated products, such as DCE and vinyl chloride, when using a palladium-on-alumina (Pd-on-Al_2O_3 or Pd/Al) catalyst. The data obtained in the experiment were fitted using a pseudo-first order model; that is, the rate of degradation of TCE, and the rate of formation of ethane, was found to be proportional to the concentration of TCE (C) in the system, provided there was excess dissolved hydrogen. Pseudo-first order degradation can be represented by the mathematical expression:

$$\frac{dC}{dt} = -kC \tag{2.1}$$

where k is a first order rate constant. In fact, k is a function of hydrogen concentration, but since hydrogen is assumed to be in excess, k may be treated as a constant. Integration of Equation 1 yields:

$$C(t) = C_0 e^{-kt} \tag{2.2}$$

where C_o represents the initial concentration of TCE.

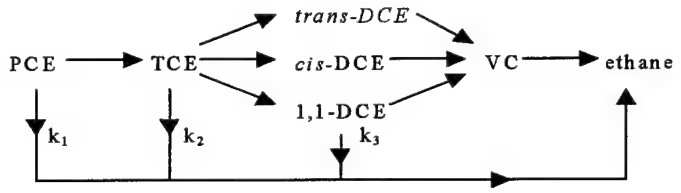

Figure 2.2. Palladium-based catalytic degradation pathways for chlorinated ethenes (Ferland, 2000).

Lowry and Reinhard (1999) also performed an experiment using metallic palladium as the catalyst to determine if rapid dechlorination was due to the palladium or if the alumina support played a role. In this experiment they found that although intermediate degradation products were detected, they never exceeded 3-4% of the initial concentration of TCE. The intermediates that were formed were also quickly degraded to ethane. From these results, Lowry and Reinhard (1999) concluded that the mechanism occurring is a complete transformation of TCE to ethane at the surface of the palladium; that is, the TCE molecules are sorbed to the catalyst surface, where the chlorine atoms are replaced with hydrogen atoms. According to Lowry and Reinhard (1999), the absence of intermediates when using Pd/Al as the catalyst may be due to the effects of the alumina support, which may be where the intermediate molecules are sorbed. This may mean that

intermediates are not detected even if present because they remain sorbed to the catalyst. They also determined that the pseudo-first order model for degradation of TCE, and formation of ethane, when using the Pd/Al catalyst, provided a good fit for the data collected. However, some deviation from first order behavior was observed; the model slightly under predicts TCE transformation at early times and over predicts it at later times. According to the authors, this may be due to a decrease in catalyst activity over the period of the experiment.

2.3.2 LAWRENCE LIVERMORE NATIONAL LABORATORY FIELD EXPERIMENT

In a field demonstration of the technology, McNab et al. (2000) constructed a dual-screened treatment well with an in-well Pd/Al reactor at the Lawrence Livermore National Laboratory (LLNL) in order to assess performance of the technology under field conditions and to identify optimal operating conditions. The site at LLNL was contaminated with different chlorinated hydrocarbons, of which TCE was the most prevalent at over 3600 ppb. Radioactive tritium was also present. The presence of tritium complicated the use of pump-and-treat technologies because of the need to dispose of the tritiated water if it was pumped aboveground. Natural attenuation of the TCE by reductive dehalogenation was not a possibility because the groundwater was aerobic. Permeable reactive barriers could not be used because the depth to the water table was 26 meters. The system used (Figure 2.3) included two catalyst beds of Pd/Al spheres, a hydrogen gas injection system, and a pumping system. Both catalyst columns were 15 cm in diameter, with lengths of 2.5 m and 2.4 m, respectively. Both columns

contained Pd/Al spheres consisting of 1% palladium metal by weight. The first column was packed with 30 kg of 0.32-cm nominal diameter spheres. The column volume was 4.418×10^{-2} m^3 and the media had a porosity of approximately 0.45. The second column was packed with 22 kg of 0.16-cm nominal diameter spheres. The column volume was 4.241×10^{-2} m^3 and the media had a porosity of approximately 0.56. According to McNab et al. (2000), the two different catalyst diameters were used as a compromise between removal efficiency, which is higher in the second column because of the greater surface area provided by smaller spheres, and the need to maintain a relatively low pressure drop across the columns. The contaminated water was pumped into the treatment well through the lower screen, where it was saturated with hydrogen gas before going through the catalyst beds. After passing through the catalyst beds, treated water was discharged into the aquifer from the upper well screen.

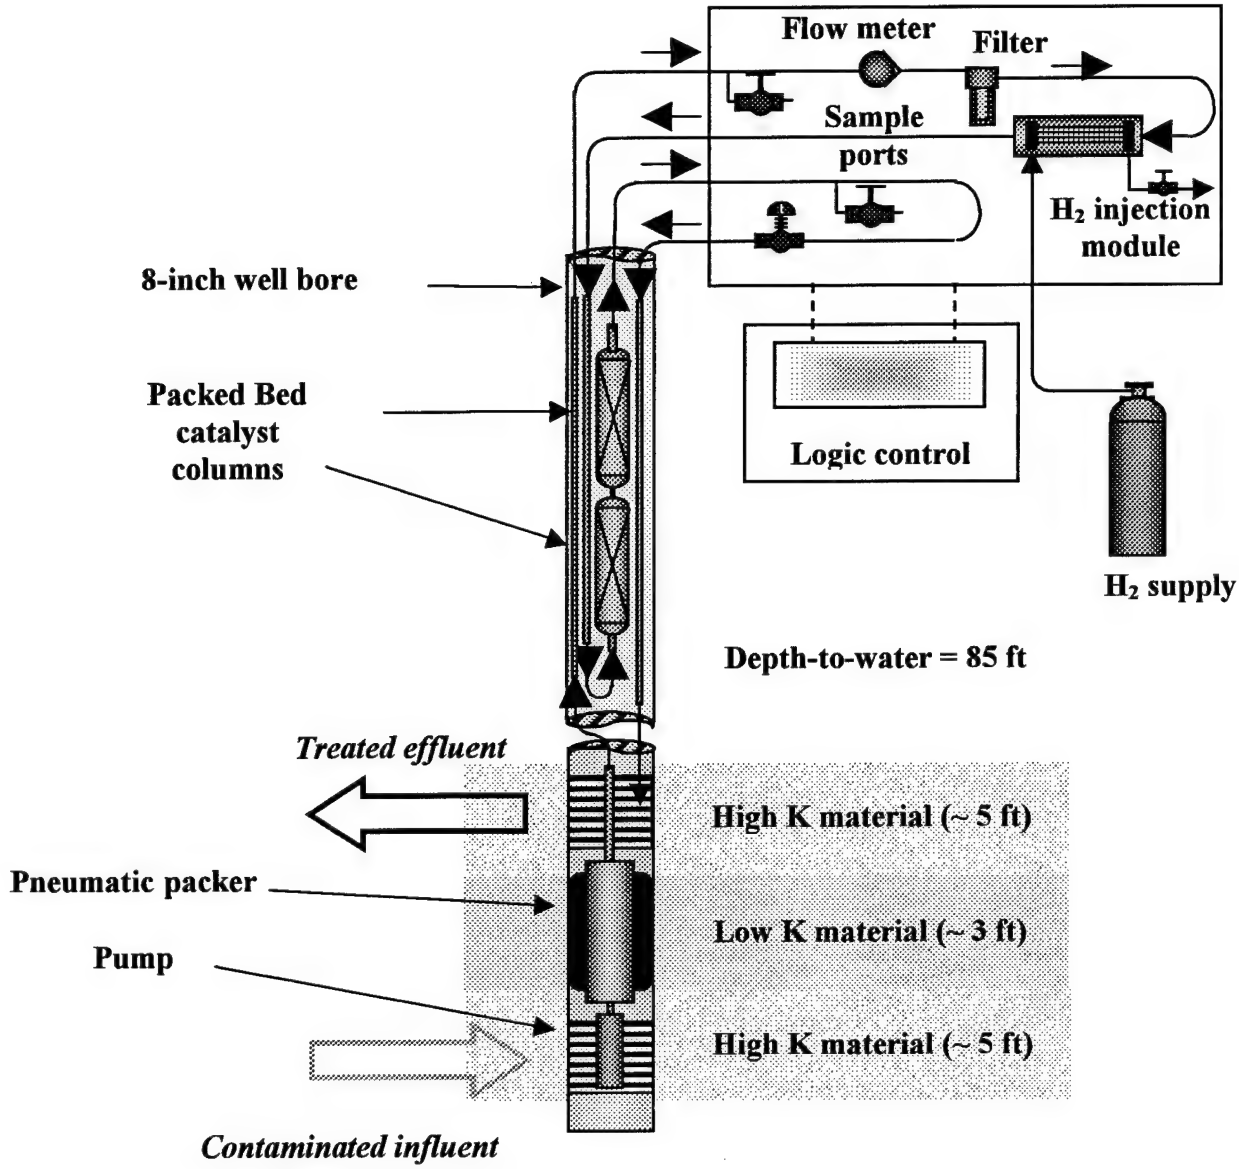

Figure 2.3. Reactive well configuration using catalytic reductive dehalogenation (McNab et al., 2000).

The field experiment was performed using two different operating cycles, four hours per day and eight hours per day. Results obtained for this field experiment were consistent with those obtained by Lowry and Reinhard (1999) in their laboratory experiments. When operating for four hours per day the system achieved over 99% removal efficiency of TCE with an effluent concentration of 0.4 to 0.8 ppb, well below

the 5 ppb limit set by the EPA, and vinyl chloride was not formed. However, when pumping for eight hours per day, the system only achieved a 93% removal efficiency with a TCE effluent concentration around 250 ppb and vinyl chloride was formed in the process, at concentrations from about 4 to 7 ppb. The system was operated at a pumping rate of 4 L/min (0.24 m^3/d), which yielded residence times of 5 and 6 minutes within the first and second catalyst columns, respectively. Hydrogen was injected at a rate of approximately 120 mL/min (7200 cm^3/d) at 3 atm pressure in order to maintain saturation.

McNab et al. (2000) suspected that the catalyst would be deactivated as time went on. McNab and Ruiz (1998) had previously determined that deactivation could be reversed by periodically removing the hydrogen supply. After the system was shut down each day, the catalyst was regenerated by purging it with three pore volumes of non-hydrogenated groundwater and then draining it to expose the catalyst to the atmosphere (McNab et al., 2000). The columns were also purged with deionized water at the end of each week and were left exposed to the atmosphere over weekends. The authors concluded that during the eight hours per day operation the palladium catalyst was deactivated through extended use, though it could be regenerated by shutting down the system. By trial and error, the authors eventually determined that operating the system between five and six hours daily produced acceptable dechlorination results.

Catalyst deactivation decreases the effectiveness of contaminant removal. Fortunately, it has been shown that catalyst deactivation during TCE dechlorination is completely reversible by different methods, such as purging with non-hydrogenated

water and atmospheric exposure, or by dosing with hypochlorite (McNab et al., 2000; Lowry and Reinhard, 2000).

2.3.4 HORIZONTAL FLOW TREATMENT WELL SYSTEM

As shown in Figure 1.1, an HFTW system consists of multiple treatment wells, some of which pump water in a downward direction alternating with wells pumping in an upward direction. Each treatment well contains a reactor that allows the contaminated water to be treated *in situ*, below the ground surface. In this study, the reactor consists of a palladium catalyst that, in the presence of dissolved hydrogen gas, degrades chlorinated contaminants. The upflow/downflow configuration allows for multiple passes of contaminated water through the reactors, permitting the attainment of very low contaminant concentrations downgradient of the treatment system (Ferland, 2000: 7). Note, to prevent short-circuiting of flow from the injection to the extraction screens of the same treatment well, HFTWs should be used either where there are two aquifers separated by an aquitard (Figure 2.4) or in an anisotropic aquifer where vertical hydraulic conductivity is much lower than horizontal hydraulic conductivity (Figure 2.5). As most aquifers exhibit hydraulic conductivity anisotropy, HFTWs should typically be appropriate for use (Domenico and Schwartz, 1998). Christ et al. (1999) developed a model to simulate groundwater flow induced by an HFTW system. In their analytical model they determine the amount of water that re-circulates through the treatment wells for specified pumping rates, well locations, and hydrologic conditions.

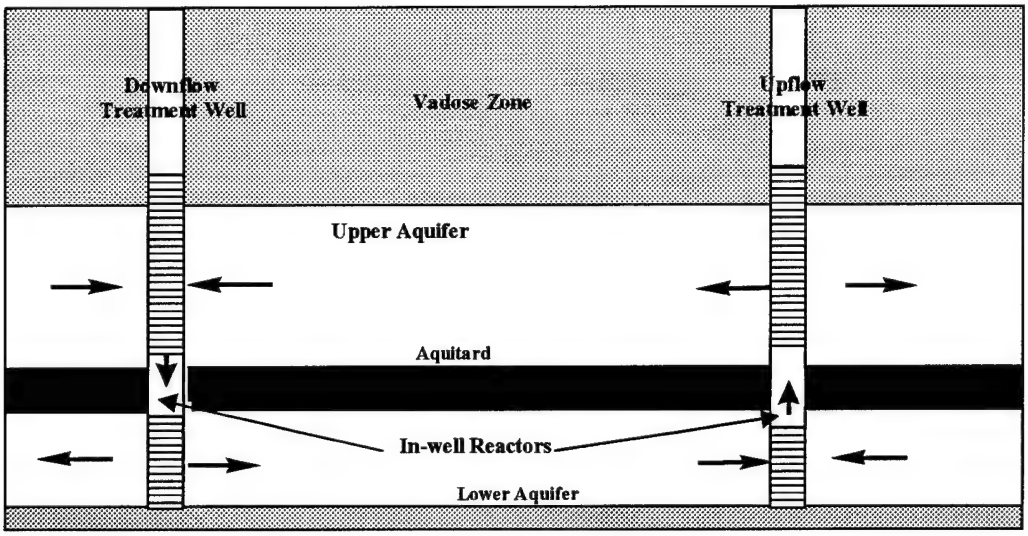

Figure 2.4: HFTW system operating in two aquifers.

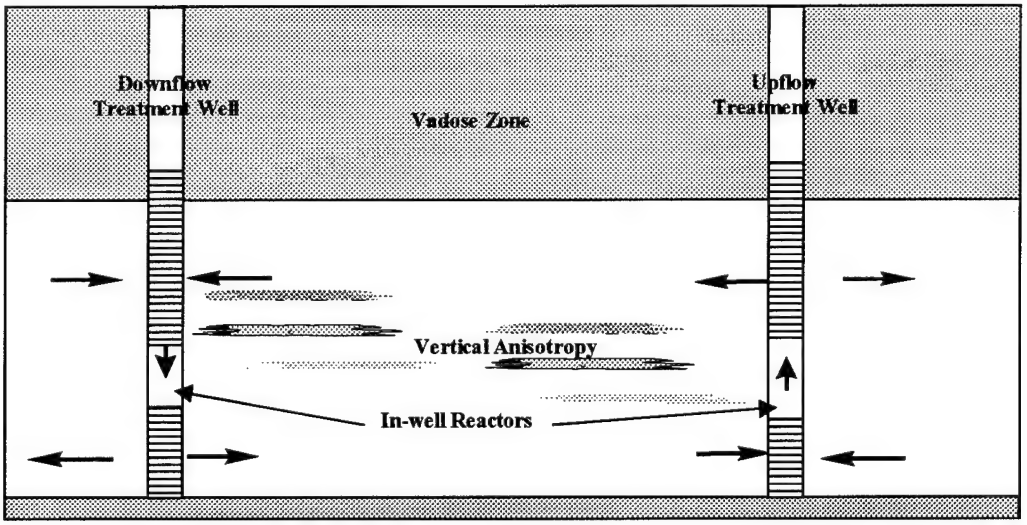

Figure 2.5: HFTW system operating in single aquifer with vertical anisotropy.

In a field study at Edwards AFB, CA, McCarty et al. (1998) demonstrated the use of an HFTW system. This site was selected for its near-ideal conditions: it contains a TCE-contaminated groundwater plume, a relatively shallow groundwater table, relatively few geological heterogeneities, and sufficient hydraulic conductivity ($>10^{-3}$ cm/s). The

site also contained two aquifers, an unconfined 8-meter thick aquifer separated from a confined 5-meter thick aquifer by a 2-meter thick aquitard (Figure 2.4). Both aquifers were contaminated with 500 to 1200 ppb of TCE. McCarty et al. (1998) evaluated *in situ* cometabolic degradation of TCE in groundwater through toluene injection, an application of bioremediation. The study was performed over a period of 410 days using an HFTW system consisting of two wells working in an upflow/downflow configuration, as described previously. The wells were spaced 10 meters apart at an angle of approximately 67.5° to the direction of regional groundwater flow (Christ et al., 1999), as depicted in Figure 2.6. Treatment occurred in bioactive zones that formed outside the injection screens of the treatment wells (Figure 2.7). In this technology, the TCE is cometabolically biodegraded in the bioactive zones. Based on results from aquifer tests and model studies, initial pumping rates for both wells were selected to be 38 L/min (54.5 m^3/d). This rate was selected because it could be obtained without excessive drawdown in the upper aquifer or pressure change in the lower aquifer. Modeling showed that for the downflow well, approximately 75% of the water pumped came from re-circulation (from the upflow well), while the other 25% came from upgradient regional flow. For the upflow well, approximately 87% of the flow came from re-circulation and the other 13% came from upgradient regional flow. At these pumping rates and distance between the wells, the model approximated a width of the plume captured of approximately 80 meters and 66 meters in the upper and lower aquifers, respectively. The model also predicted a removal of approximately 83% of TCE with each pass through a bioactive zone, and an estimated 95-97% overall removal, comparing upgradient and downgradient concentrations. The reason overall removal rate was greater than single-pass removal

was because of re-circulation of treated water in the HFTW system. In the final stage of the actual experiment, the flow rates were set to 25 L/min (36 m^3/d) in each well instead of 38 L/min. Using mass balance, the experimenters were able to determine that the actual recycle rate was 91.5% for the downflow well and 84% for the upflow well. TCE removal efficiency was determined to be 86% for a single pass and approximately 97-98% for the system overall. This study demonstrated that implementation of an HFTW system can significantly increase the overall efficiency of a groundwater treatment system by re-circulating contaminated water.

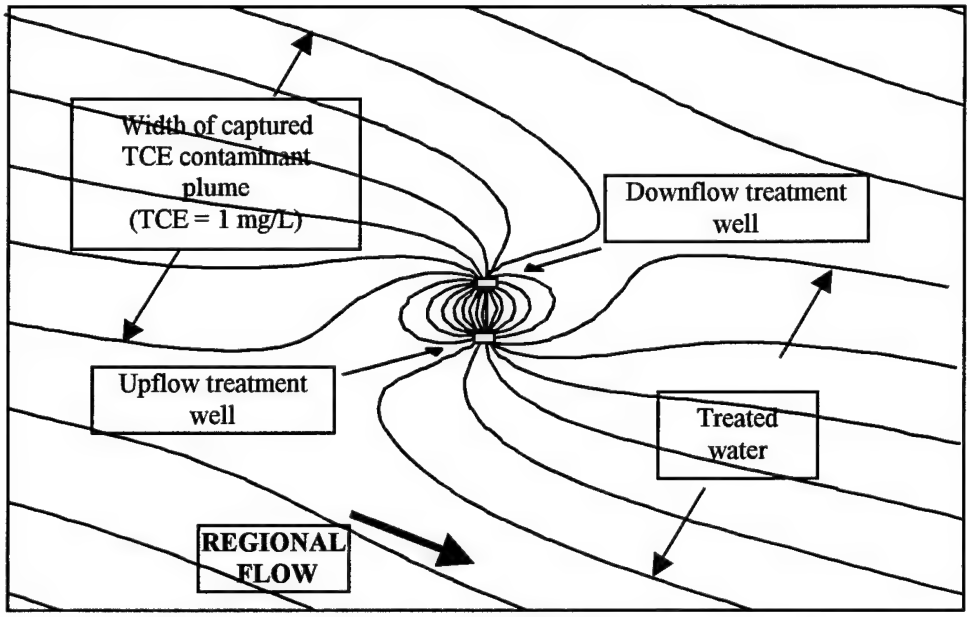

Figure 2.6. Plan view of HFTW system showing re-circulation (McCarty et al., 1998).

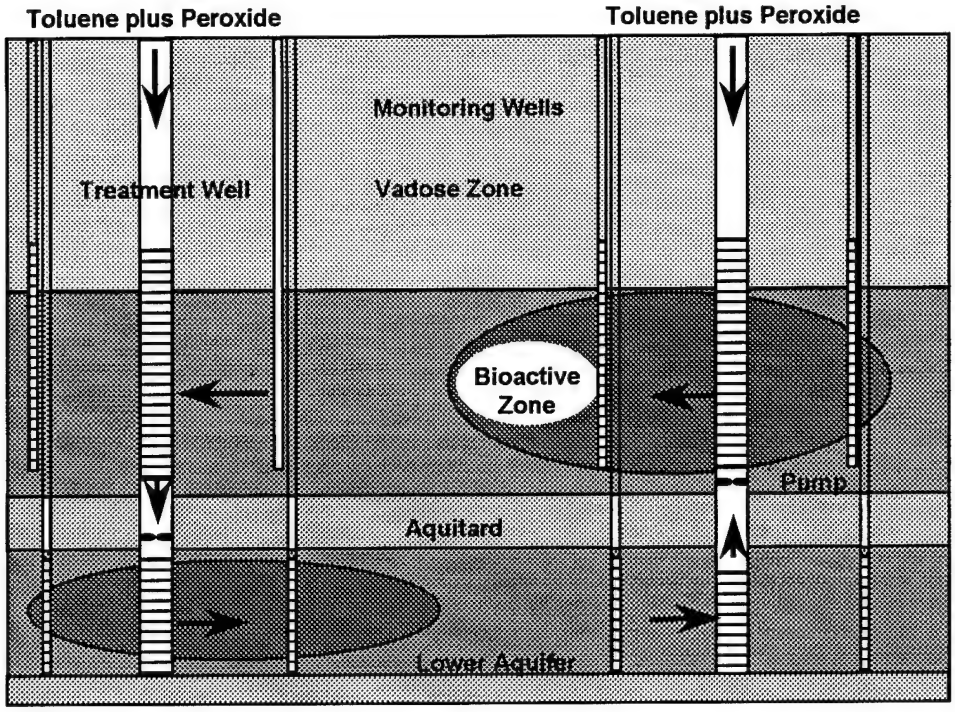

Figure 2.7. Cross-section of cometabolic TCE degradation HFTW system (McCarty et al., 1998).

2.4 TECHNOLOGY MODELING

Ferland (2000) developed an analytical model to aid in the design of HFTW systems with in-well palladium reactors, and Huang and Goltz (1998) developed a numerical model to describe contaminant transport in the subsurface.

2.4.1 ANALYTICAL MODEL OF HFTW SYSTEM WITH IN-WELL PALLADIUM REACTORS

Ferland (2000) developed an analytical model to simulate first-order degradation of TCE using an in-well palladium reactor in a two-well HFTW system. He assumed a TCE half-life degradation constant, k, based on the work of McNab et al. (2000).

Ferland's (2000) primary treatment objective was to achieve an overall degradation efficiency high enough so the downgradient concentration of TCE was below 5 ppb, the drinking water MCL. The secondary objective of the effort was to ensure the capture of the contaminated groundwater plume.

Using Ferland's (2000) model, a project manager interested in implementing the technology may determine an in-well reactor volume, the required distance between treatment wells, and pumping rates, given a specified capture zone width and upgradient contaminant concentration. To apply the model, certain environmental parameters such as the regional groundwater Darcy velocity, the thickness of the aquifer, and the angle of flow of the regional groundwater relative to a line connecting the HFTW treatment wells must be known or assumed. Using Ferland's model, various combinations of reactor volume, distance between treatment wells, and treatment well pumping rates that meet treatment objectives could be determined. By comparing the various combinations, the project manager could decide how the technology could be cost-effectively implemented at a site.

Note that the analytical model developed in Ferland (2000) is intended for use as a screening model. Due to the numerous simplifying assumptions incorporated into the model (e.g. aquifer homogeneity, steady-state flow and transport), use of the model to actually design a remediation system is limited. To better understand and simulate the technology under real conditions, numerical models, which do not make all the simplifying assumptions of analytical models, are needed.

2.4.2 NUMERICAL MODELING

A number of numerical models have been developed to simulate groundwater flow and contaminant fate and transport. One such model was developed by Huang and Goltz (1998). The model is the only three-dimensional model that has been applied thus far to simulate the operation of an HFTW system (Garrett, 1999). MODFLOW (Harbaugh and McDonald, 1996) is used to simulate groundwater flow in the aquifer, while a fate and transport model, based on the transport model MT3D, is used to simulate contaminant (TCE) fate and transport in the aquifer. The three-dimensional model uses finite differences, and solves the partial differential equations describing TCE fate and transport using a self-adaptive, partial-implicit approach (Garrett, 1999). Huang and Goltz (1998) developed this model to simulate the *in situ* bioremediation technology implemented by McCarty et al. at Edwards AFB (see section Section 2.3.4).

In the model, Visual MODFLOW is used to generate a finite difference grid (Figure 2.8). Well locations and pumping rates, initial conditions, and boundary conditions are specified. Hydraulic conductivity for each cell in the grid must also be specified. MODFLOW permits specification of hydraulic conductivity anisotropy, which, as noted earlier, is important when using an HFTW system. With this information, MODFLOW can calculate steady state hydraulic heads and velocity fields. Then, the velocity fields, along with contaminant initial concentrations and boundary conditions, can be used by the fate and transport model to determine how contaminant concentration varies over space and time. The model simulates physical and chemical transport processes such as advection, dispersion, and sorption. It may also be modified to simulate different types of treatment processes at the treatment wells. The model has

been used to simulate aerobic cometabolism (Garrett, 1999), and in this work can be modified to model first-order degradation through palladium-catalyzed dehalogenation.

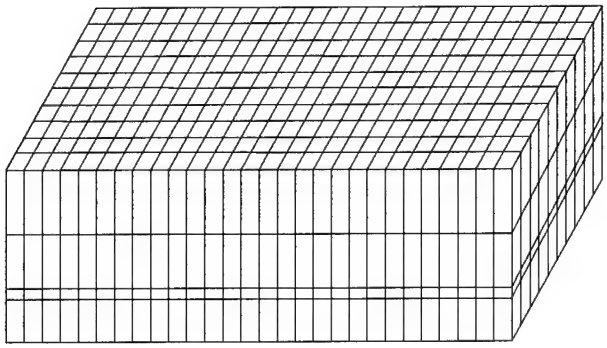

Figure 2.8. Example 3-D finite difference grid (Garrett, 1999).

2.4.3 MODEL PARAMETERS

Note that in both the analytical and numerical models described above, there are two types of parameters. One type, which we will call environmental parameters, includes such things as regional groundwater flow velocity and direction, hydraulic conductivity, contaminant concentrations, etc. The other type of parameters, which we will refer to as engineered parameters, includes such things as location and pumping rates of wells, and treatment technology specifications (e.g. reactor volumes and hydrogen injection rates). Pumping rate and reactor size determine the residence time of the contaminated water in the reactor. Larger reactor residence times allow for increased TCE removal during a single pass through the reactor. Larger residence times can be accomplished by increasing the size of the reactor or by decreasing the pumping rate. However, changing either of these two parameters could have significant effects on cost and/or overall efficiency for the system. An increase in flow rate decreases single-pass

efficiency but increases re-circulation. On the other hand, decreasing pumping rates would increase single-pass efficiency but decrease re-circulation. Also, increasing pumping rates increases operating costs, and increasing the size of the catalyst column would increase capital costs. Obviously, there are optimal sets of parameter values that achieve the required system performance most cost effectively. In the next section we will look at optimization methods that may be used to determine the "best" engineered parameters to achieve some treatment objective.

2.5 OPTIMIZATION

According to *Merriam Webster's Collegiate Dictionary* (1993), optimization is "the mathematical procedure involved in making something, such as a design, system, or decision, as fully effective as possible." Optimization usually involves determining a maximum or minimum of an objective function that is bounded by some constraints. A simple example of an optimization problem is as follows:

Maximize: $f(x, y) = x + y$
Subject To: $0 \leq x \leq 10$
$0 \leq y \leq 12$

In this case the solution is evident: the maximum value of $f(x, y)$ occurs when $x = 10$ and $y = 12$, which yields $f(x, y) = 10 + 12 = 22$. Based on this example, useful terminology may be defined (Garrett, 1999: 2-2).

Objective - A goal toward which effort is directed. In this case the goal is to get the highest possible value of $f(x, y)$ given certain limitations, which are the constraints.

Objective function - The mathematical representation of the objective. In this case it is:

Maximize: $f(x, y) = x + y$

Decision Variables or Parameters - Variables included in the objective function, in this case x and y.

Constraints - Limits on the values that may be assigned to the decision variables or on other aspects of the system, such as on the objective function or a combination of the decision variables.

Subject to: $0 \leq x \leq 10,\ 0 \leq y \leq 12,\ x - y \leq 0$

Solution - The values assigned to the decision variables that yield the optimal value for the objective function. In this case

$x = 10, y = 12$

Candidate Solution - A set of values assigned to the decision variables that may or may not be a good solution.

Evaluation Function - A function used to determine if candidate solutions provide "good" values. In this case it is the same as the objective function.

Real world problems may or may not be trivial. In any event, they must be studied carefully in order to ensure complete understanding of the problem at hand. The decision variables, the objective function, and the constraints may not initially be well defined, and great care must be exercised in defining them to ensure the final solution to the problem is useful and meaningful.

Every problem has a solution space, or landscape, that may be obtained by plotting the objective function value for each possible set of parameters. Smooth landscapes with one or few optimal values (Figure 2.9) are generally easy to solve using simple optimization techniques while rugged landscapes with multiple local optimal

values (Figure 2.10) are usually very difficult to solve. A hill climbing algorithm is a search method that looks for the best solution in an area, or neighborhood, by following a path of the highest gradients, or changes in terrain, it can find. This type of algorithm would easily solve a problem with a smooth landscape. However, it would almost certainly get "stuck" at a local optimum, which would most likely not be the global optimum, in a rugged landscape. To deal with these difficult problems, heuristic search methods have proven useful.

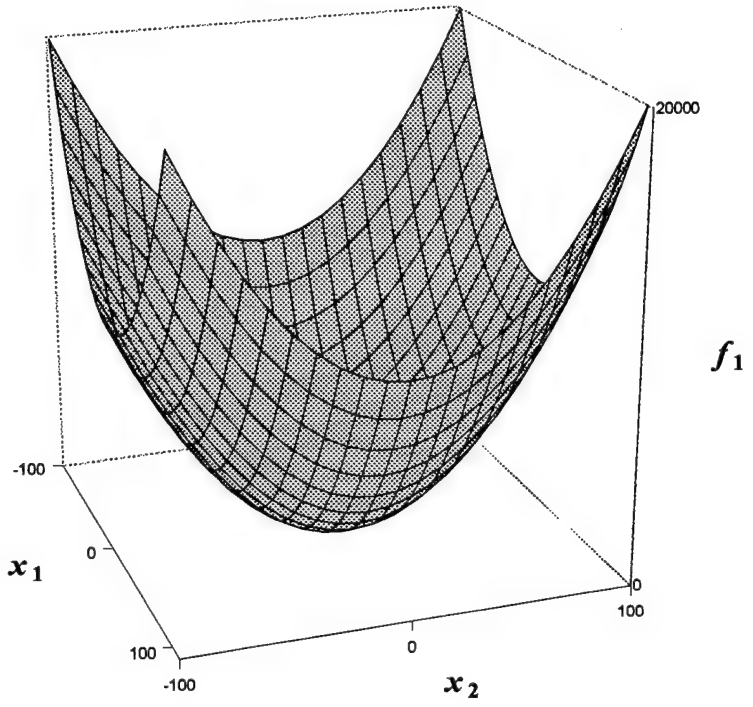

$$f_1(x) = \sum_{i=1}^{n} x_i^2$$

Figure 2.9. Example of *smooth* landscape. f_1 for $n=2$; $-100 \leq x_i \leq 100$ (Garrett, 1999).

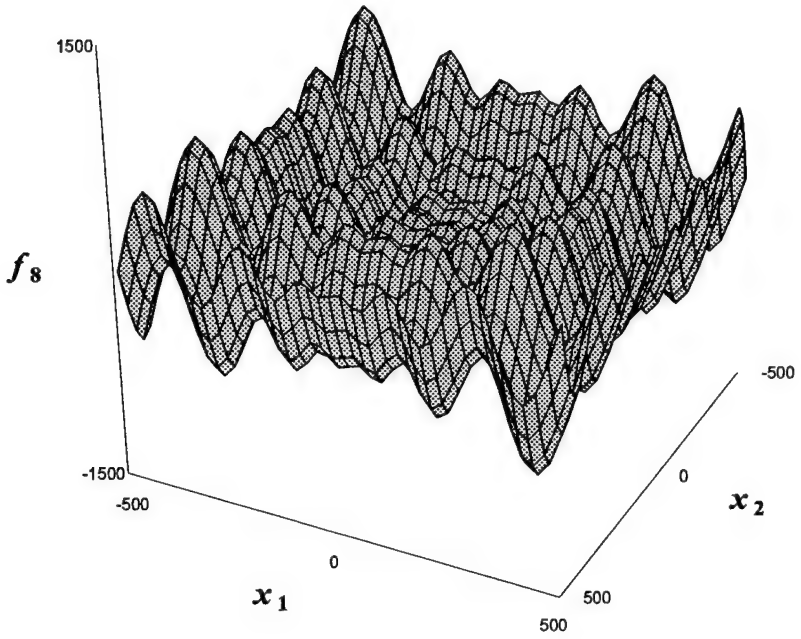

$$f_8(x) = -\sum_{i=1}^{n} x_i \sin(\sqrt{|x_i|})$$

Figure 2.10. Example of *rugged* landscape. f_8 for $n=2$; $-500 \leq x_i \leq 500$ (Garrett, 1999).

2.6 MODERN HEURISTIC SEARCH METHODS

In this section, an introduction to different types of heuristic search methods, or heuristics, is presented. *Heuristics* are "techniques which seek good (i.e., near optimal) solutions at a reasonable computational cost without being able to guarantee either feasibility or optimality, or even in many cases to state how close to optimality a particular feasible solution is" (Reeves and Beasley, 1992: 6). Heuristics are used to solve complex optimization problems and they tend to be simpler to understand than other methods. A category of algorithms that belongs in this category is evolutionary

algorithms (EAs) and a specific type of EA is genetic algorithms (GAs). The discussion focuses on EAs and specifically on GAs.

2.6.1 INTRODUCTION

When presented a complex problem, such as the optimization of a groundwater remediation technology, different techniques may be used to solve it. These techniques range from using intuition and experience to using linear or non-linear programming, random search methods, calculus-based search methods, evolutionary algorithms, etc. (Garrett, 1999: 2-14; Goldberg, 1989: 2). Although more than one technique may be chosen to solve a given type of problem, some techniques are faster while others may yield better solutions. The method used to solve a problem depends on the characteristics of the problem. Some problems may be easily formulated mathematically, which would allow for their solution using relatively simple techniques such as linear optimization. Other problems may require an unreasonable amount of time or computing resources to solve using these techniques. These problems may be solved using modern heuristic search methods.

Heuristics are typically classified into several broad categories such as greedy construction methods, neighborhood search routines (NSR), relaxation techniques, partial enumeration, decomposition and partition approaches, and others (Reeves and Beasley, 1995: 12). Some of the most common heuristic techniques can be categorized in one of the above categories and include: simulated annealing, a type of NSR; tabu search, another type of NSR; and EAs, a type of random search. These methods were developed as attempts to simulate naturally occurring processes, which make them fairly easy to

understand. For example, simulated annealing is analogous to the thermodynamic processes of heating and cooling steel, tabu search was motivated by attempts to imitate intelligent processes, specifically memory, and EAs were modeled after genetic structures. We will focus here on EAs, and specifically GAs, primarily because they have been successfully applied in the past to solve very similar problems (Garrett, 1999). Other reasons for our focus are the ready availability and ease of use of commercial GA packages.

2.6.2 EVOLUTIONARY ALGORITHMS

EAs are heuristic search methods inspired by biological processes of natural selection and evolution. These algorithms include at least the following: genetic algorithms, evolution strategies, evolutionary programming, and genetic programming, all of which are similar but possess differing characteristics and strengths (Bäck, 2000: 59, Dasgupta and Michalewicz, 1997: 3).

Over the years, creatures must evolve in order to adapt to changing environments and to survive. The existence of species, and of specific members within a species, is dictated by survival of the fittest. In general, stronger and smarter creatures are able to outcompete weaker ones for food and shelter, therefore they survive. This allows strong individuals to mate and pass on their genetic information to subsequent generations. By combining strong individuals, species become stronger and are able to better adapt to their environment by continuing to improve over time. EAs are based on the principle that the biological processes that result in evolution can be modeled and used to solve complex optimization problems. Most evolutionary algorithms maintain a population of

individuals, or *chromosomes*, that evolves over time (Garrett, 1999). Each individual represents a candidate solution that can be evaluated for *fitness*, which is a representation of how well that individual meets the objective function. The fitness of each individual is compared to that of other individuals in the population, and those individuals with better fitness are selected to survive and mate with each other. Mating of individuals produces a new generation of individuals that is, hopefully, more "fit" than its parent generation. The new individuals are then evaluated for fitness, and the fitter ones are selected for survival. The process repeats until time runs out or until the population converges on a single solution (Garrett, 1999: 2-22 - 2-23). The general algorithm for an EA is as follows (Bäck, 2000: 61):

initialize the parent population;
evaluate each individual in the parent population for fitness;
while termination condition is not true **loop**
    recombine the parent population to create an offspring population;
    mutate the offspring population;
    evaluate each individual in the offspring population for fitness;
    select individuals for survival to become new parent population;
    determine termination condition;
end loop

As the general algorithm above indicates, most EAs are comprised of four major operations: recombination or reproduction, mutation, evaluation, and selection (Bäck, 2000: 61). In general, the recombination operator takes two or more individuals from the parent population and combines them to produce one or more children. The mutation operator randomly changes part of a chromosome. The evaluation operator evaluates each individual in the population and assigns them a fitness value. And, the selection operator determines which individuals survive to pass on their genetic information to

later generations. These operators allow the population to improve while exploring the landscape of the problem (Garrett, 1999).

In general, individuals in EAs are represented as a coded vector, which is their genetic information, or genotype. This genetic information can be decoded to provide the physical characteristics of the individual, which include values for each variable associated with it, as well as permitting calculation of fitness. These "physical" characteristics are known as phenotype (Fogel, 2000: 23-24). During recombination, individuals' genotypes are combined to form new phenotypes, which determines the phenotype of the new individuals.

All types of evolutionary algorithms function on the basic use of the four operators described above. The primary differences between different types of EAs are how they represent individuals and how new individuals are generated through recombination and mutation. The basic principles of each are (Bäck, 2000: 60):

1) **Genetic algorithms**. Use recombination, or crossover, as the main operator, while mutations are a "background operator" that is used rarely on a probabilistic basis. Selection is also a probabilistic operator and binary representation of individuals, or candidate solutions, is most common, although real-valued representation is also used. Population size in different generations is usually identical.

2) **Evolution strategies**. Mutations and recombinations are the essential operators and mutations are normally distributed. Selection is deterministic. Individuals are represented using real-valued vectors. The population size typically differs between generations.

3) **Evolutionary programming**. Emphasizes mutations and does not incorporate the recombination of individuals. Mutations are also normally distributed. Selection is probabilistic. Mainly used with real-valued vectors, but originally developed to evolve finite-state machines, or mathematical logic.

4) **Genetic Programming**. Used to automatically develop computer programs. Each individual represents a complete computer program in a suitable programming language (Bäck and Fogel, 2000: xxvii).

Genetic algorithms have been chosen for this effort because of their simplicity and availability. GAs have been used widely to solve complex engineering and science problems, including optimization of groundwater remediation systems. In a similar undertaking, Garrett (1999) developed a genetic algorithm to optimize the HFTW system tested by McCarty et al. (1998) (Section 2.3.4).

2.6.3 GENETIC ALGORITHMS

Genetic algorithms (GAs) were first developed by John Holland at the University of Michigan in the late 1960s and 1970s (Holland, 1975). More recently, David Goldberg published a book on GAs that has become widely accepted (Goldberg, 1989). GAs have numerous advantages. Firstly, GAs are very robust; that is, they are capable of solving a wide array of problems and have been used in many fields, including groundwater remediation (Ritzel et al., 1994; Cieniawski et al., 1995; McKinney and Lin, 1994; Guan and Aral, 1999; Rogers et al., 1995; etc.). GAs are also very flexible. Since variables are typically encoded, usually in binary strings (explained in the following section), they can represent many different types of functions. As an example, in some

groundwater remediation optimization problems, a bit in a binary string is selected to represent whether a pumping well is turned on or off. A string of binary bits may also represent a real-valued number such as pumping rate or angle of flow. Because this representation of variables is usually coded in binary, operators can easily manipulate the variables (explained in the following sections). Also, this type of representation and the simplicity of the operators make GAs non-problem specific, which means that the same GA code, with minimal changes, may be used to solve many different types of problems. Since GAs rely on random searches, while still exploiting historical information, they can be expected to find a globally optimum problem solution (Goldberg, 1989: 1-2). The details of how they do this, using chromosomes and evolutionary operators, follows.

2.6.3.1 CHROMOSOME REPRESENTATION

In order to better understand how GAs work, an introduction to GA terminology is necessary. In GAs, each individual is known as a chromosome and is usually coded. A chromosome is the genotype that is manipulated by the operators. Evaluation decodes each chromosome into a phenotype in order to calculate a fitness value. The genotype is the coded space while the phenotype is in the actual space. For example, $X_1 = 8$ is a phenotype value, while (1000) is the genotype for X_1. A chromosome consists of a number of parameter values, known as *genes*. For example, the values for pumping rate and distance between wells would each be a gene. The possible values of each gene, or each position in the gene, are known as *alleles*, and the position of a variable in a gene is called its *locus* (Eshelman, 2000: 64; Reeves, 1995: 153).

In order for GAs to operate on candidate solutions, the solution must be converted into a phenotype that can be manipulated and evaluated by the algorithm, while at the same time remaining understandable by the user. In general, GAs code the parameters that comprise a solution set. The most common way to code the parameters of a candidate solution is with a binary representation (Goldberg, 1989: 7). Binary representation consists of zeroes and ones to represent decimal numbers. The previous example may be used to demonstrate how binary coding works:

Maximize: $f(x, y) = x + y$
Subject To: $0 \leq x \leq 10$
$0 \leq y \leq 12$
x, y must be integers

Binary coding for x and y is very simple, with each variable represented by four digits. The maximum value for x, 10, would be represented as the binary number 1010, and for y, 12, would be represented as the binary number 1100. Thus, the chromosome that represents the candidate solution $x = 10$, $y = 12$ could be represented as 1010 1100, or as 10101100, with the four leftmost digits representing x and the four rightmost digits representing y. However, in GAs this representation is not important as long as the evaluator can decode it. In this case, the eight bits could be arranged in any order, however, the representation presented above is the most common because it allows the user to keep better track of each gene in the chromosome. Binary representation is very common in GAs because of its simplicity and because of the ease with which the operators can manipulate binary numbers (Garrett, 1999: 2-25).

2.6.3.2 OPERATORS

As in most EAs, GAs use the four basic operators of crossover, mutation, evaluation, and selection. The following discussion includes an overview of each of these operators along with some basic examples. The examples will use binary chromosome representation for simplicity.

2.6.3.2.1 CROSSOVER

Crossover is the main operator used in GAs; it is the primary means of generating new individuals. The crossover operator exchanges sections of "parent" chromosomes. Crossover is analogous to the idea of reproduction in biological processes, where chromosomes from two individuals are combined to form one or more individuals.

Since individuals are chosen based on probability, as explained later in the section describing the selection operator, a crossover operator is most likely to combine chromosomes from two individuals with relatively high fitness values. By choosing the fittest individuals from a population to create new ones, each new generation should inherit the best characteristics of the previous generation, while minimizing the number of bad characteristics.

Crossover operations can be done in several ways, one of the most common being one-point, or simple, crossover. One-point crossover involves taking two individuals from the parent population and, randomly, selecting a crossover point. As illustrated in Figure 2.11, the first offspring (Child X) is composed of the portion of Parent A to the left of the crossover point and the portion of Parent B to the right of the crossover point. Similarly, Child Y is composed of the portion to the right of the crossover point in Parent

A and the portion on the left of the crossover point in the Parent B. In this example the crossover point is after the fifth bit (Garrett, 1999).

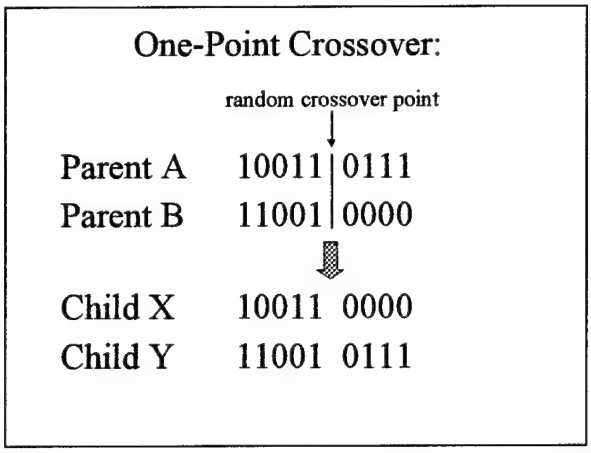

Figure 2.11. One-point crossover (Garrett, 1999).

The principles used for simple crossover can be extended to what is known as multiple-point crossover. This type of crossover is similar to simple crossover, but more than one point is chosen for crossovers.

Another type of crossover is the uniform crossover. It involves two parents and generates one or two children. First, a binary string that acts as a pattern, or bit-mask, is randomly generated. This pattern is used to determine which child gets information from which parent. In the loci where the pattern's locus is a '1', the child gets the alleles from the first parent, and in the loci where the pattern is a '0', the child gets the alleles from the second parent as illustrated in Figure 2.12. If a second child is generated, the reverse is true (Garrett, 1999).

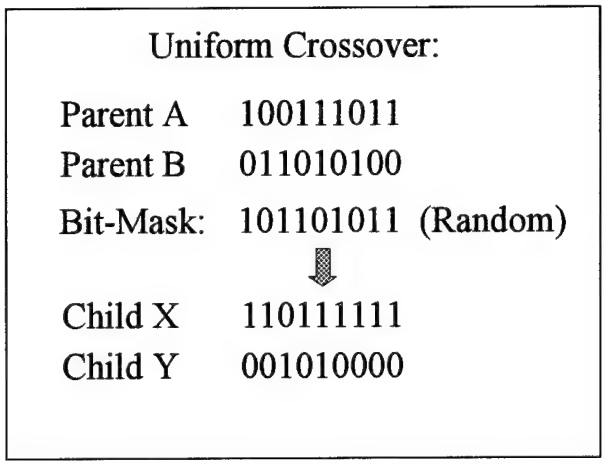

Figure 2.12. Uniform crossover (Garrett, 1999).

2.6.3.2.2 MUTATION

Crossovers explore areas of the landscape that are included in the original populations. Mutation is an operation that is used in GAs mainly because some good genetic material may never be explored through crossover only (Goldberg, 1989: 14). The mutation operation involves a random modification of a chromosome and it occurs after a new generation is generated through crossovers. After mutations are performed on a population, the new individuals formed replace those that were mutated. In our example, or any problem represented in binary form, mutation involves changing a single allele at a random locus in a chromosome from 1 to 0 or from 0 to 1. The rate at which mutations occurs is usually set to a very low value (on the order of 0.01) because too many mutations may destroy good genetic information, causing the algorithm to essentially become a random search. However, the performance of a GA has been shown to be enhanced by a low rate of mutations (Garrett, 1999: 2-29). Over the years, researchers have developed different types of mutations (Bäck et al., 2000: 238). Two types are creep mutation and jump mutation. Jump mutations are as described above, a

change of an allele at a random locus on a random individual. Creep mutations involve increasing or decreasing the value of a random gene by one unit. As an example of creep mutation, the chromosome 10101010 would be mutated to 10101011 (increasing) or to 10101000 (decreasing).

2.6.3.2.3 EVALUATION

The evaluation operator uses the fitness function to evaluate each individual in a population to determine how well the individual meets the objective function. Each individual is assigned a fitness value from the evaluation, which is later used for selection. According to Cieniawski et al. (1995: 402), the evaluation procedure is typically very time consuming, because there are so many evaluations that must be performed, one for each member of the population. Also, when dealing with real-world problems, each evaluation may involve complex and involved calculations or simulations (Garrett, 1999: 2-30). One method that may be used to decrease the amount of time required to run a GA is the use of parallelization, which involves using multiple computers to evaluate multiple individuals concurrently.

2.6.3.2.4 SELECTION

Once a new generation has been created from the parent generation and each individual has been assigned a fitness value, the selection operator is used to determine which individuals will be selected to survive into the subsequent generation. The most common method is to select a mix of members from the parent and children populations (Garrett, 1999: 2-30 – 2-32).

Selection may be performed either through stochastic or deterministic methods. In a simple GA, the process is stochastic in nature, with the individuals having the highest fitness function values being assigned higher probabilities of survival. Another common method is the elitist model, a deterministic method in which the members with the highest fitness values are automatically chosen to survive and pass their genetic material on to subsequent generations (Reeves, 1995: 166). Tournament style selection is performed by randomly selecting a group of individuals and comparing their fitness. The individuals with the highest fitness from each group are selected to reproduce. These tournaments are repeated until enough individuals are chosen.

After these four operations are completed and a new population is chosen, the process begins all over again, until a stopping criterion is met. This stopping criterion is usually based either on time, number of generations, or on the measured improvement in the solutions; that is, the difference between the Nth and the (N + 1)th generations in satisfying the objective function falls in some pre-described range. In general a GA will need a good number of generations to ensure that good solutions are found.

2.7 GAs IN GROUNDWATER REMEDIATION

GAs and other optimization techniques have been extensively used to optimize groundwater monitoring and remediation technologies (Rogers et al., 1995; Ritzel et al., 1994; Cieniawski et al., 1995; McKinney and Lin, 1994; Guan and Aral, 1999; Garrett, 1999; Culver and Shoemaker, 1993; Dougherty and Marryott, 1991; Rogers and Dowla, 1994; Yoon and Shoemaker, 1999). Most of these optimization efforts have been geared towards pump-and-treat systems or water quality monitoring systems.

Rogers et al. (1995) used artificial neural networks (ANNs) and a GA to minimize the cost of operating an existing pump-and-treat facility with 28 injection or extraction wells. In this endeavor their goal was to determine the cheapest way to contain the contaminant plume and maximize the amount of contaminant mass removed from the aquifer. They trained the ANN to recognize patterns for good and bad solutions from a groundwater flow and contaminant transport algorithm in order to significantly lower the running time of the algorithm. A binary GA was used in which each available pumping well was represented by a single bit, which represented on or off. In the end, they were able to find good solutions in a reasonable amount of time. In similar undertakings Aly and Peralta (1999) and Rogers and Dowla (1994) have also used ANNs and a GA to optimize a pump-and-treat groundwater cleanup system.

Ritzel et al. (1994) used two GAs to optimize reliability and cost of a hydraulic containment, or pump-and-treat, system. With this algorithm they determined how many injection and extraction wells to install, where to install them, and how much to pump from each. In this problem they gave 16 locations for possible wells and allowed the GA to find the optimal solution. They used binary coding, with each chromosome being a representation of the pumping (extraction or injection) rates for all of the wells. Then they compared the solutions obtained using the GAs and the mixed integer chance constrained programming (MICCP) method, which had been used previously by Morgan et al. (1993). According to Morgan et al. (1993), the MICCP method they used found solutions that were very close to optimality. The GA used by Ritzel et al. (1994) found solutions that were just as good as the ones found by MICCP in a shorter period of time.

McKinney and Lin (1994), Guan and Aral (1999), and Vasquez et al. (2000) have also used GAs in other pump-and-treat optimization efforts.

Numerous other approaches have been attempted in efforts to optimize pump-and-treat systems. These include: nonlinear programming (Ahlfeld et al., 1988; McKinney and Lin, 1994), quasi-Newton methods (Culver and Shoemaker, 1993), simulated annealing (Marryott et al., 1995), and combinations of these and others.

In a similar effort, Garrett (1999) developed a real-valued parallel GA to optimize the operation of a two-well HFTW system established to promote *in situ* aerobic cometabolic bioremediation of TCE-contaminated groundwater. This system is the same one that McCarty et al. (1998) demonstrated at Edwards AFB. The objective of the optimization effort was to minimize the cost of the system while achieving a downgradient concentration below the MCL for TCE of 5 ppb. In this system a number of parameters were of interest because in order for cometabolic TCE degradation to occur there must be plenty of dissolved oxygen in the groundwater, microorganisms, and a food source for the microorganisms, which was toluene in this case. The parameters optimized were: number of treatment wells, treatment well locations, treatment well pumping rates, amount of oxygen added to the water, amount of toluene added to the water, oxygen injection schedule, and toluene injection schedule (Garrett, 1999).

3.0 METHODOLOGY

3.1 OVERVIEW

In this chapter, the application of a GA to optimize palladium-catalyzed *in situ* dechlorination of TCE at a specific site is discussed. In the first section, a discussion of our initial problem formulation, including the objective, cost, and fitness functions, is presented. In the second section, the model selected to simulate contaminant fate and transport and assumptions on aquifer characteristics are discussed. In the third section, the genetic algorithm (GA) selected for this endeavor and its operators are discussed. In the fourth section a description of how the GA will be applied to the problem in the optimization effort is presented. Finally, in the fifth section, a revised fitness function, based on the cost of destroying contaminant mass, is presented. This revised fitness function was employed to help deal with artificial results that were obtained from using the GA with the originally proposed fitness function. This study is very similar to Garrett's (1999) study in that both involve application of horizontal flow treatment wells (HFTWs) to attain a specified remedial objective. Because of these similarities, this effort will closely resemble Garrett's (1999). The studies differ in that Garrett (1999) optimized a model of a bioremediation technology, *in situ* aerobic cometabolism, while we are optimizing a model of a chemical process, palladium-catalyzed dehalogenation. In the current study, the single-pass treatment efficiency of the palladium reactor is inversely related to well flow rate (since lower flow rates result in higher residence times in the reactor), while with the bioremediation technology single-pass treatment efficiency was independent of flow rate (Mandalas et al., 1998).

3.2 PROBLEM FORMULATION

In order to optimize this technology, we must formulate an objective function that accurately represents the goals of the technology application and a fitness function that effectively allows the GA to determine the best solutions. This section presents a description of the initial fitness function developed, an explanation of its shortfalls, and a description of the fitness function that was ultimately used.

3.2.1 OBJECTIVE FUNCTION

As described in Section 2.5, the objective function is a mathematical representation of a goal towards which an effort is directed. In this case the effort is directed at cost effective containment and remediation of the contaminant plume. As stated in previous chapters, the EPA has set a TCE drinking water maximum contaminant level (MCL) of 5 ppb, or 5 μg/L (Masters, 1997). Ideally, the remediation system would be installed and operated to capture the entire plume of contaminant and to clean up the contaminated groundwater to achieve a downgradient concentration below the MCL for the least possible cost.

The following objective statement is a formulation of the above-stated goals:

Implement palladium-catalyzed in situ destruction of TCE-contaminated groundwater using HFTWs, at minimum cost, so as to completely capture a plume of TCE-contaminated groundwater so that the concentrations downgradient of the treatment system are within regulatory limits.

Following the format of the example provided in Section 2.5, we can define the objective function as follows:

Minimize: Cost (3.1)
Subject To: capture of entire TCE plume
downgradient concentrations of TCE < 5 μg/L
specified constraints on decision parameters

Cost is the total cost in dollars of installing and operating the system. The downgradient concentrations are measured at various observation wells downgradient of the treatment system. The decision parameter constraints define the minimum and maximum allowable values for the engineered parameters being optimized. These constraints will be discussed and specified in Sections 3.2.2 and 3.4.

3.2.2 DECISION PARAMETERS

In order to implement this technology at a contaminated site, a project manager must know which parameters are under his/her control. Likewise, he or she must determine which parameters may be changed in order to optimize the system. As described in Section 2.4.3, there are two different types of parameters associated with a groundwater remediation model, environmental and engineered parameters. In this section we will focus on the engineered parameters.

The following engineered parameters are under the remediation manager's control and may be changed to affect the operation of the system:

(1) Number of treatment wells

(2) Location of treatment wells

(3) Pumping rate of treatment wells

(4) Volume of catalytic reactor

When formulating the solution to the problem, these variables must be defined mathematically and constrained. As an example, if the number of treatment wells is not constrained, an infinite number of wells and an infinite number of well locations would be possible, making the problem far too difficult to formulate and to solve. In order to formulate a tractable problem and impose constraints, we must have an idea of the situation in which we are employing the treatment system.

In this study we are implementing the treatment system in a homogeneous aquifer similar to the one at Edwards AFB, CA, which is described in Section 2.3.4 (McCarty et al. 1998; Garrett, 1999). Garrett (1999) also applied a technology based on the Edwards AFB aquifer, therefore some of the parameter ranges we define will be based on Garrett (1999).

As previously explained, the number of treatment wells in the system must be constrained in order to have a tractable problem. In this study, we will constrain the number of wells to two since we are attempting to determine the optimal configuration and pumping rates for the wells. Therefore, the number of treatment wells will not be considered a decision parameter.

The location of the treatment wells is a parameter that is defined by two sub-parameters, horizontal distance between the wells and the angle of the wells relative to regional groundwater flow direction. The location of the wells is a very important parameter that can have significant effects on system efficiency. Changing the distance between the wells or the angle of flow affects the interflow between the wells, thereby affecting the overall efficiency of the system.

As discussed in Section 2.4.3, well pumping rate is another parameter that greatly affects the efficiency of the system because it affects interflow between the wells and residence time in the catalytic reactor. Therefore, it affects both single-pass efficiency and overall efficiency of the system. Since we are dealing with an HFTW system, in which treatment wells move water from one aquifer to another, pumping rates should be the same for both wells in order to maintain an overall water balance between the aquifers. Otherwise, decreases or increases in pressure in the confined aquifer and cones of depression or mounding in the water table aquifer may develop.

The volume of the catalytic reactor is also a very important parameter, because it, along with pumping rates, determines the residence time of the contaminated water in the reactor, thus determining single-pass treatment efficiency. In this study the diameter of the column will be kept constant at 20 cm, but its length will be allowed to change. The volume of the reactor also has significant effects on the cost of the system, as we will see in the following section.

The decision parameters that will be used in this study are: 1) distance between treatment wells, or well separation distance (meters); 2) angle of a line connecting the two treatment wells relative to the regional groundwater flow direction, or angle of flow (degrees); 3) pumping rate (meters^3/day); and (4) length of the catalytic reactor (meters). All of these parameters must also have upper and lower bounds specified. We shall wait to define these until after we have discussed the model and other assumptions that may affect these constraints.

3.2.3 COST FUNCTION

In this section we focus on the costs associated with installing and operating an HFTW system with in-well palladium reactors. Christ (1997) developed annualized cost equations for cleaning up TCE-contaminated groundwater using *in situ* cometabolic bioremediation, which also uses HFTWs. Garrett (1999) used these equations to develop his cost function in an HFTW system. Because of the numerous similarities between the systems, these formulas can easily be applied to our system.

As explained by Christ (1997), there are two types of costs associated with groundwater cleanup systems, capital costs and operating costs. Capital costs are those costs associated with the initial purchase, installation, and start-up of the system, annualized over the expected life of the project, which we will assume to be 20 years. In this case capital costs include the wells, pumps, the Pd/Al catalyst, and any other equipment that may be needed to start-up the system. Capital costs can be defined as (Christ, 1997: 3-31; Garrett, 1999:3-8):

$$\text{TCC} = Cost_C \times N \times f(y) \tag{3.2}$$

where

TCC = Total capital costs annualized over y years [$/year]

Cost_C = The capital/installation/initial start-up costs per well [$]

N = Number of treatment wells

$f(y)$ = A function used to annualize the capital costs

$$= \frac{i \times (1+i)^y}{(1+i)^y - 1}$$

i = The expected average annual interest rate = 6%

y = The duration of the project = 20 years

As previously stated, we will use two wells in this optimization problem. The cost of installation of a well is approximately $10,000 (Mandalas, 1997:C-38)) plus the cost of the catalyst, which is approximately $270/kg (McNab et al., 2000). A column of packed, 1.6-mm nominal diameter beads of Pd/Al has an approximate density of 1.2 g/cm^3 (McNab et al., 2000), and an effective porosity of 0.42 (Lowry and Reinhard, 2000). This means that the catalyst in a 4-m long, 20-cm diameter column would cost approximately $23,600. Annualized over 20 years, at an interest rate of 6%, the TCC of such an HFTW system would be $5860/year.

Operating costs are costs associated with daily operation of the system and consist of pumping costs and other costs of running the system such as the cost of catalyst regeneration. In this study, regeneration costs are assumed to be insignificant compared to other operating costs, and will be neglected. Therefore, operating costs are defined as the yearly cost associated with pumping groundwater and can be mathematically expressed as (Christ, 1997; Garrett, 1999):

$$\text{TPC} = \frac{TE \times N \times Cost_E}{\eta} \tag{3.3}$$

$$\text{TE} = \frac{P \times 1}{1000} * (hrs/day) * (days/yr)$$

$$\text{P} = \gamma \times H \times Q$$

where

TPC = Total annual pumping costs [$/yr]

TE = Total annual energy used by pumping [kW-hr/well-year].

P = Power used to lift water [kg-m^2/sec^3 or Watts]

γ = Specific gravity of water [$kg/m^2\text{-}s^2$] = 9.807 m/s^2 * 1000 kg/m^3

= 9807 $kg/m^2\text{-}s^2$

H = The effective distance that the water must be lifted [m]

Q = Flow rate in each well [m^3/s]

N = Number of wells

$Cost_E$ = Cost of energy [$ / kW-hr]

η = Wire-to-water efficiency (accounts for electric power lost because not all electricity is converted to pumping power)

For our study, the distance the water must be lifted is assumed to be 15m, which is approximately the distance from the bottom of the confined aquifer to the top of the water table at the Edwards AFB site. To this we must add the headlosses through the system, the main one being the headloss through the catalyst column. Other headlosses in the system due to piping, etc. are assumed small in comparison with the loss due to flow through the catalytic reactor. In order to calculate headloss (hdls) through the packed-bed catalytic reactor column of length L, we use the Hazen equation (Metcalf & Eddy, 1991:268):

$$hdls = \frac{1}{C} \times \frac{60}{10+T} \times \frac{L}{d_{10}^{\ 2}} \times V_h \tag{3.4}$$

$$V_h = \frac{Q}{A}$$

where

C = Coefficient of compactness = 1200 for a packed column

T = Temperature of groundwater in °F = 65°F

L = Length of the catalyst column [m]

d_{10} = Effective grain diameter size [mm] = 1.6 mm

V_h = Superficial filtration velocity (Darcy velocity) [m/d]

Q = Pumping rate [m^3/day]

Area = Cross-sectional area of the catalyst column [m^2]

The cost of energy we will use is $0.196/kW-hr (Mandalas, 1997), and we will assume a value for η of 0.6. This wire-to-water efficiency incorporates, and is the product of, thermal efficiency (E_m) and pump efficiency (E_p). E_m typically ranges from 0.8 to 0.95 for electric motors and E_p typically ranges from 0.6 to 0.85 (Bouwer, 1978:184-187). In this problem we are also assuming that the system will operate 24 hours per day, 365 days per year.

To get an idea of what the operating costs would be for a typical system, we can use the information above. Given a pumping rate of 50 m^3/d with a 4 m catalyst column we would expect headlosses through the column to be approximately 1.66 meters. Therefore, the power used to lift water through the column and the well would be approximately 94.54 Watts, which translates to 828.17 kW-hr/well-year of energy. At the given energy cost, operation of a two-well system would cost approximately $540/year.

The equations developed for capital costs and operating cost can be combined to determine the total annual cost (TC) of the system, which is the cost we want to minimize.

$$TC = TCC + TPC \tag{3.5}$$

Using the values in our example, we can show that the cost of a two well HFTW system with in-well palladium reactors would be approximately $6400/yr.

Since our objective for this problem is to minimize the cost of the system, we can now develop an objective function:

minimize:	TC	(3.6)
subject to:	downgradient TCE concentrations ≤ 5 ppb	
	capture of contaminant plume	
	specified constraints on decision parameters	

3.2.4 FITNESS FUNCTION

A fitness function is the metric that the GA uses to determine which solutions are superior to others. Hence, the fitness function must be an accurate representation of how well each individual meets the objective function previously formulated. Total cost is measured using the equations previously formulated, while downgradient TCE concentration and capture of the contaminant plume are calculated by the fate and transport model, as described in Section 3.3. However, since we will be measuring whether or not a plume is captured simply by measuring downgradient concentrations (that is, if the plume is not captured, downgradient concentrations will clearly exceed regulatory limits), we will ignore the constraint and include it as part of the downgradient concentration constraint. We must now develop a fitness function that represents how well these objectives are met.

GAs, unlike other types of optimization techniques, such as linear programming, are difficult to constrain numerically. Therefore, we may develop a fitness function that incorporates the constraints on the problem by placing weights on parameters to indicate their importance (Garrett, 1999). Because we are trying to minimize our cost while achieving downgradient TCE concentrations below a given value, we can incorporate the

concentration constraint in the fitness function. In this case, concentrations should be penalized based upon how much they exceed the desired value. Concentrations at or below the desired value should be given a value of zero because they meet the requirement. Note that there is no "reward" for being below the desired value. In order to combine downgradient concentrations and cost into a fitness function, we need to normalize both functions so we can compare them on the same scale.

As previously mentioned, one of our objectives is to achieve TCE concentrations of 5 ppb (5×10^{-6} g/L), the regulatory limit, downgradient of the treatment wells. For this constraint, we must decide upon a proper range of normalized values. As was discussed, a concentration that meets the regulatory limit should have the same effect on the fitness function, no matter how far below the limit this solution is. However, if the concentration is above the limit, we must be able to differentiate between one that barely fails the standards, such as 10 ppb, versus one that is not anywhere near the desired concentration, such as 1 part per million (ppm or 10^{-3} g/L). Garrett (1999) defined a normalized penalty function ($TCE_{adjusted}$) for violating the downgradient concentration constraint as follows:

$$TCE_{adjusted} = 0 \qquad \text{for } TCE_{recorded} \le TCE_{desired}$$

$$TCE_{adjusted} = \frac{TCE_{recorded}}{TCE_{desired}} - 1 \qquad \text{for } TCE_{recorded} > TCE_{desired} \qquad (3.7)$$

where

$TCE_{recorded}$ = Max TCE concentrations downgradient of the treatment system

$TCE_{desired}$ = The regulatory limit, or 5×10^{-6} g/L

In the same manner we can normalize cost. In the previous section we calculated the annualized cost of a system as approximately $6000. We will assume this approximates the cost of a typical system. Therefore, we can normalize cost as follows:

$$Cost_{adjusted} = \frac{Cost_{actual}}{Cost_{desired}} \tag{3.8}$$

where

$Cost_{adjusted}$ = The normalized annual cost of a specified system for use in the fitness function

$Cost_{actual}$ = The annualized cost of a specified system

$Cost_{desired}$ = The yearly cost of a typical system, or $6000/yr

Now that the cost and concentration functions have been scaled, we can add them together to calculate a fitness function that accounts for both. Therefore our fitness function is:

$$fitness = Cost_{adjusted} + TCE_{adjusted} \tag{3.9}$$

Looking at the cost equations (Equation 3.5), we would expect the cost portion of the fitness function would be relatively smooth and, therefore, easy for the GA to optimize. However, the downgradient concentration of TCE is determined by very complex relationships, which are calculated using the fate and transport model. We would expect the function would be rugged, requiring more time for the GA to optimize. If both cost and downgradient concentration are weighted the same through the entire optimization process, the GA will try to optimize cost, since it is easier. This could cause the GA to quickly minimize cost, without attaining the concentration goals (Garrett, 1999). Hence, we must devise a solution search strategy that does not allow the GA to concentrate too heavily on the cost, to prevent the GA from converging on a cost

effective solution that does not meet the concentration requirements. Therefore, a fitness function that attempts to prevent this from happening must be developed. Garrett (1999) observed similar behavior in his study, which led him to develop a time-based fitness function that initially more heavily weights solutions that meet downgradient regulatory constraints and subsequently concentrates on lowering system cost. This function is defined as follows:

$$fitness_{time} = Cost_{adjusted} \times time + TCE_{adjusted} \tag{3.10}$$

where

$$time = \frac{CurrentGenerationNumber}{MaximumGenerationNumber}$$

Hence, if the optimization consists of 100 generations, during the first generation, cost will be weighted 1/100th of the weight given to downgradient concentrations. As the GA approaches the final generations, cost and concentrations approach equal weights.

3.3 NUMERICAL MODEL

One of the most important steps of the modeling process is mapping the site onto a model grid. The importance of this lies in that it will determine not only the accuracy of the model, but also the computational time required, which is very important when using GAs because of the number of iterations required. In order to apply the model to a real site we must convert site characteristics into model parameters. First, the site must be discretized into a three-dimensional finite difference grid using Visual Modflow. Visual Modflow has a windows-based interface where the user may specify aquifer characteristics, initial conditions, boundary conditions, and other modeling parameters.

These parameters include location and concentration of contaminant sources, aquifer hydraulic conductivity, hydraulic gradient, size of the area being modeled, and grid size of the area. An important feature of the program is that it allows the user to specify non-uniform grid sizes (Figures 3.1), which allow for a very large reduction in computational run-time of the model. Visual Modflow may also be used to place injection and/or extraction wells in the aquifer model. The files generated by this program are saved in a format that can be read by the FORTRAN program MODFLOW.

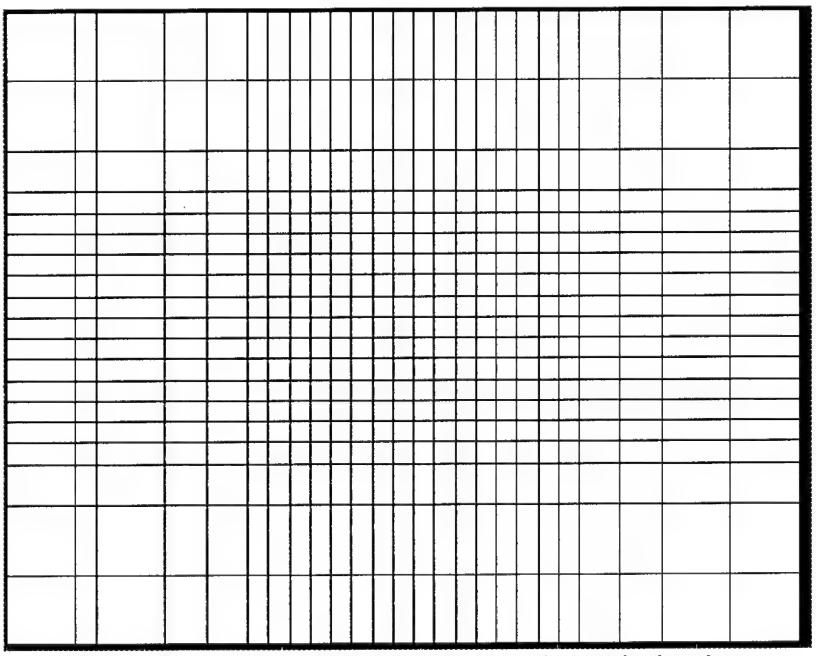

Figure 3.1. Non-uniform grid used in optimization.

MODFLOW is a program whose source code is freely available (Harbaugh and McDonald, 1996). The program reads the files that are generated by Visual Modflow and uses them to calculate groundwater flow conditions in the aquifer. MODFLOW is capable of computing transient flow or steady-state flow conditions. In this study we are assuming steady flow conditions. MODFLOW also generates output files, which along with the parameters generated by the GA (well location, catalyst column length, etc.) are used by the fate and transport model, which is based on MT3D, to calculate

concentrations over time throughout the aquifer. The fate and transport model calculates these concentrations using the flow fields generated by MODFLOW and by simulating advection and dispersion of the contaminant through the aquifer. A reaction component has been added to MT3D to allow for simulation of different treatment technologies, such as aerobic cometabolic bioremediation (Garrett, 1999) and first order decay in a catalytic reactor, as used in this study.

As previously stated, the contaminant fate and transport model determines the concentrations for every cell in the grid through every time step of a simulation. It does this by modeling advection, dispersion, and sorption of the contaminant. In our model, we conservatively assume the effects of sorption are negligible.

To summarize the workings of the model, initially a site model must be developed using Visual Modflow. Then, the GA generates individuals (solution sets), each consisting of a set of engineered parameter values. The GA calls on MODFLOW to calculate the flow fields and the fate and transport model to calculate the contaminant concentrations. These concentrations are then returned to the GA in order to calculate the fitness value of the individual solutions.

3.3.1 MAPPING A HYPOTHETICAL SITE ONTO A MODEL GRID

As previously mentioned, the mapping of a hypothetical site onto a model grid is a very important process; therefore, we spend more time discussing the details of this effort. The site layout and characteristics used are based on geological conditions at Edwards AFB, CA as described in Section 2.3.4. Garrett (1999) also used a model based on this site. Four layers were used in the model. The top layer is 8 meters deep and

represents the vadose zone and part of the saturated zone. It is at this level that we set the water table for the aquifer. The second layer represents an unconfined aquifer 8 meters deep. The third layer is a low permeability layer, or an aquitard, 2 meter deep that separates the second and fourth layers. The fourth layer is a confined aquifer that is 6 meters deep.

Initially, the size of the site to be modeled was chosen to be a 100 meter by 100 meter square. It was divided into a grid of 40 columns and 40 rows of equal size, which totaled 1600 grid cells in each layer. However, the computation time for such a grid was on the order of approximately 45 minutes on an 800 MHz computer for one simulation of the numerical model. Since it would take over 1 day to run one generation of 50 individuals, it was infeasible to run the GA at this resolution. Further refinement of the grid changed the grid to a size of 115 meters in length and 91 meters in width divided into 25 columns and 19 rows. The cell size ranged from 3m x 3m to 10m x 10m, depending on location (Figure 3.2). The higher resolution cells (smaller grid sizes) are in the areas near the pumping wells while the lower resolution cells are at the boundaries of the grid (Figure 3.2). This configuration contained approximately 30% of the number of cells as the original, and at this resolution one run of the model took approximately 5 minutes, which means that approximately five generations of individuals could be run daily.

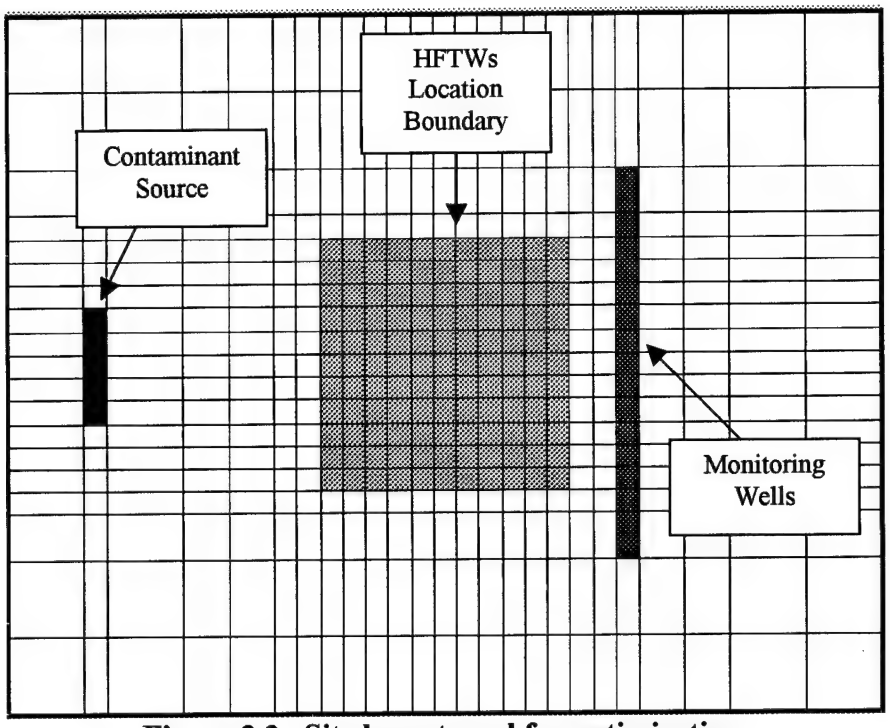

Figure 3.2. Site layout used for optimization.

Groundwater in the model was simulated to flow from left to right using a hydraulic gradient of approximately 0.0043. The hydraulic conductivity was set to 2.95 m/d in all directions in both aquifers, making this a model of a homogeneous, isotropic aquifer. A contamination source area was located near the leftmost boundary of the grid and rows of observation wells were placed in layers 2 and 4 near the rightmost boundary of the grid. The length of the source was set at approximately 15 m long and one cell wide (3 m). This source area, through dispersion, generated contaminant plumes approximately 25 to 35 meters wide. The contaminant concentration was specified at 5 mg/L, which was greater than the 1 mg/L concentration found at Edwards AFB (McCarty et al., 1998). The higher concentration was used to make it more difficult for the technology to meet downgradient requirements. At this source concentration, the downgradient concentration objectives could be attained with some of the well configurations and pumping rates tested, but not with all. A full-scale implementation of

this model could, depending on the width of the plume, use more than two wells. However, the current study is constrained to look at two-well solutions.

3.4 GENETIC ALGORITHM SOFTWARE SELECTION

A FORTRAN 77-coded GA was selected for this application because of its availability and ease of use. The GA was developed by David L. Carroll from the Department of Aeronautical and Astronautical Engineering at the University of Illinois at Urbana-Champaign. It is the only readily available FORTRAN-coded GA, and it can be downloaded from the World Wide Web (http://www.staff.uiuc.edu/~carroll/ga.html). Different versions of this GA have been used previously in optimization of problems involving lasers (Carroll, 1996) and in medical prostate implant research (Yang et al., 1998). In this study, Version 1.7.0 of the GA, which is the latest free available version, was used.

The GA is binary-coded and it is set up to maximize a function. Since our objective is to minimize a function, we multiplied the fitness function (Equation 3.9) by negative one, which essentially converts the GA to a minimization algorithm. Adapting the GA to our problem was straightforward. A general function is included with the GA so any user can test the GA before applying it to his or her problem. To adapt our problem to the GA, we replaced this general function with our simulation model and the cost equations previously formulated. For each individual solution, the simulation model returned values for concentrations at the observations wells, which, along with the calculated cost, were used by the GA to determine the fitness of the individual.

The GA allows us to specify an initial population of individuals rather than having the GA randomly generate an initial population. Fifty individuals were initially created using a systematic combination of five values for each of the four engineered parameters (pumping rate, distance between wells, angle of flow, and length of catalyst column). Combinations of an extreme low (e.g., Q = 9.41 m^3/d), a middle low (e.g., Q = 26.23 m^3/d), a middle value (e.g., Q = 45.85 m^3/d), a middle high (e.g., Q = 64.07 m^3/d), and an extreme high (e.g., Q = 81.59 m^3/d) were chosen (see Appendix A). By doing this, we attempted to ensure the GA searched the entire solution space and did not converge on a local optimum.

Either single point crossover or uniform crossover can be selected. Because we generated the initial population to cover the entire solution space, single point crossover was used. As suggested in the GA documentation (Carroll, 1998), a crossover rate of 0.7 was selected, which means that on the average 70% of all chromosomes, or individuals, are combined in each generation. Two children were generated in each crossover, which means that both parents are replaced by the children produced.

The GA is also coded to allow two types of mutations, jump and creep (described in Section 2.6.3.2.2). As suggested by Carroll, the rate of jump mutations was set to $1/n_{pop}$, or 0.02, and the rate of creep mutations was set to $2/n_{pop}$, or 0.04.

The GA driver is set for tournament selection, which is performed by randomly selecting a group of individuals from the population and allowing the most fit individuals from each group to mate (Carroll, 1996). This process continues until the initial size of the population is reached. Elitism, or the automatic reproduction of the chromosome

with the best fitness in each generation, was also used. This function prevents the loss of good chromosome strings (Carroll, 1996).

Finally, as previously stated, a population of 50 individuals was originally generated. The population of 50 remained constant through all generations. The number of generations was set to 65. This value was chosen as a compromise between allowing the GA enough generations to converge and keeping computing time manageable. Garrett (1999) ran his GA for 100 generations. However, the best fitness did not significantly improve after approximately 60 generations and after approximately 70-75 generations the chromosomes appeared to converge. Running the GA with a population of 50 for 65 generations involves over 3000 model simulations, which, if each simulation averages 5 minutes, would take approximately 11-12 days to finish.

The GA also requires that a range be set for each parameter. The range of the values selected were as follows:

x_1: Q - Pumping Rate [m^3/d]	$1 \leq x_1 \leq 90$
x_2: d - Horizontal Well Separation [m]	$6 \leq x_2 \leq 30$
x_3: α - Angle of Flow [degrees]	$0 \leq x_3 \leq 360$
x_4: L - Length of Catalyst Column [m]	$1 \leq x_4 \leq 15$

The lower limit for pumping rate was selected because the model does not run at a pumping rate of zero. In addition, a pumping rate of zero, or even a very small pumping rate, would not be very effective in capturing the contaminant plume. The upper limit was selected based on Garrett's (1999) work. When the pumping rate is too high, it causes excessive drawdown, which may cause the grid cell where the well is located to dry up and report a contaminant concentration of 0.0 mg/L in the cell (Garrett, 1999).

Comparisons of model simulations using pumping rates above 90 m^3/d and below 90 m^3/d did not show improvements in concentrations for the higher rates. The bounds for the second parameter, d, were selected solely based on characteristics of the finite difference grid used to model the site. The lower limit was selected to ensure that the two wells were not placed in contiguous cells. The upper limit was selected to ensure the wells were placed in the area of high resolution in the center of the grid. The angle of flow was selected to range across a full rotation to ensure that the best well orientation was selected. Finally, since we desired the technology to be applied *in situ*, the length of the catalyst column was constrained by the distance from the bottom of the confined aquifer to the top of the unconfined aquifer, which is 15 meters.

Another option that may be used with this GA is the use of a micro-GA (μ-GA). The μ-GA included with Carroll's software is a type of GA that uses a small population and, by using elitism, searches for an optimal solution. If a population of five individuals is chosen, the μ-GA initially generates five solutions randomly, then it evaluates each of them and, using crossovers and mutations, generates new individuals. At the beginning of each new generation the five individuals are compared to each other to check for convergence, and if the individuals are too similar to each other, the best individual is kept and the other four are randomly re-generated. Therefore, this approach is like a combination of random search with GA. Uniform crossovers using a rate of 0.5 were used with the μ-GA because uniform crossover is more random, and in this case we are trying to determine an optimal solution through a pseudo-random search rather than by converging to a given solution. In our problem, the μ-GA results were compared with the

results of the regular GA. The same model was used in each μ-GA optimization as in the GA optimization.

3.5 APPLICATION OF GA TO PROBLEM

The GA was applied using our fitness function (Equation 3.10). As previously described, the groundwater model returns concentration values for the observation wells, of which the highest is used to calculate the normalized penalty function for concentration. This penalty function, along with the cost, is used to calculate individual fitness.

As previously shown in Figure 3.2, the center of the grid was selected as a high-resolution area in which the treatment wells could be placed. Since placement of the wells is dependent on two parameters, well separation distance and angle of flow, after the GA generated an individual parameter set, the grid location of the wells had to be calculated based on the size of the cells. The center point of the grid was used to determine treatment well location, using the specified values of d and α. Observation wells were placed downgradient of the treatment wells. A total of ten observation wells were placed in the second and fourth layers. The highest concentration observed in the ten wells was taken to be the concentration downgradient.

The duration of each simulation was set to 600 days. Tests of the model showed that after 600 days there is not much change in concentrations through the aquifer. However, for methodology verification, each optimal, or near optimal, solution will be tested for a simulated duration of 1000 days. During verification, we will also use a

higher resolution (38 x 29 - Figure 3.3) grid to check the accuracy of the solution that was obtained based on use of a coarser grid.

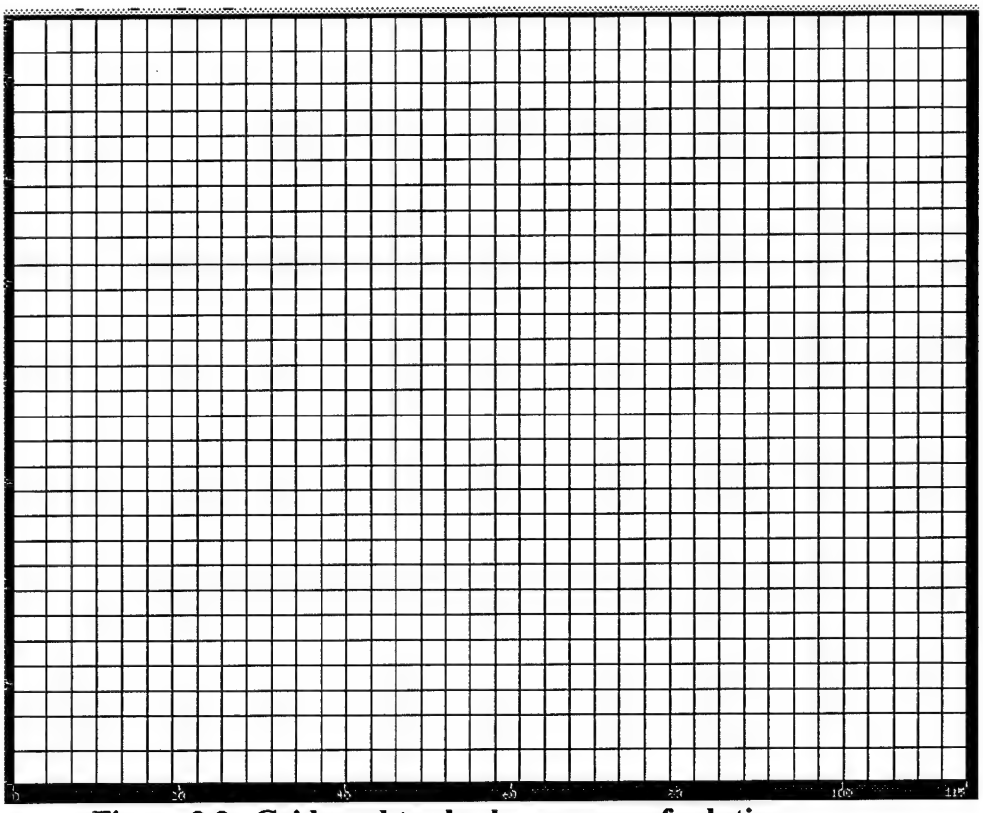

Figure 3.3. Grid used to check accuracy of solution.

3.6 REVISED FITNESS FUNCTION

The initial run of the GA, using the methodology previously described, generated unexpected results. By the final generation, the GA had converged on solutions containing similar values for all the parameters. The optimal pumping rate (Q) converged to approximately 15-20 m^3/d, the distance between the wells (d) to between 25-30 m, the angle of flow (α) to 180 degrees (downflow well closest to the source), and the length of the catalyst column (L) to 3-5 m. The best solution had the following values: $Q = 15\ m^3/d$, $d = 30$ m, $\alpha = 180^\circ$, and $L = 3.0$ m. Unfortunately, based on these

results, it appears that the GA is converging on an artificial solution which, although minimizing the objective function, is not realistic. The μ-GA optimization resulted in a similar finding; an optimal value for α was found to be 0° (upflow well closest to the source). Note that with a value of α of 180°, all the treated water discharged by the downflow well into the lower aquifer flows to the upflow well (interflow = 100%) [similarly, a value of α of 0° indicates all the treated water discharged by the upflow well into the upper aquifer flows to the downflow well]. α = 180° means that contaminated water from the upper aquifer that is captured by the downflow well is treated twice (once in the downflow well, a second time in the upflow well), but the contaminant plume in the lower aquifer is not treated at all; rather it is spread out around the re-circulation zone between the two treatment wells [α = 0° has a similar effect in the upper aquifer]. Because of the location of the observation wells downgradient of the treatment wells, and the simulation time of 600 days, the lower aquifer plume flowing around the re-circulation zone did not reach the monitoring wells over the course of the simulation. In order to solve this problem, we attempted placing monitoring wells cross gradient of the treatment wells. However, even though these monitoring wells detected higher concentrations than the wells downgradient, concentrations were still not high enough to force the GA into realistic solutions. The reason for this appears to be that the contaminant plume did not reach steady state after 600 days because more contaminant was entering the system than was being degraded. A solution to this problem could be to add more wells to the system, however we are constrained to two wells in this study. Running each simulation for longer periods of time (1000-1500 days or more), would allow the higher concentrations of the contaminant plume in the bottom aquifer to reach

the observation points, thereby generating worse fitness values for those solutions. However, these longer simulations could not be accomplished as the computer resources available for this study would require two to three weeks to run such long simulations.

Because of these problems, a new fitness function was developed with the same objective of optimizing system performance and cost. As previously, an objective statement, from which we'll develop an objective function and a fitness function, must be defined. A measure that can be used to account for the performance of the system is the amount of contaminant mass destroyed. However, the cost of the system must also be taken into account. A system using 15-m catalytic columns would degrade more TCE than one using 5-m columns, but it would also be much more expensive. Therefore, the following objective statement represents the stated goals: *implement palladium-catalyzed in situ destruction of TCE-contaminated groundwater using HFTWs to destroy the maximum amount of TCE possible at minimum cost.* As the statement indicates, this is a problem of both maximization (mass destroyed) and minimization (cost). To implement these objectives in a GA we can combine these two goals into one that reflects both of them. The objective chosen was to minimize the cost per TCE mass destroyed.

Following the format of our previous objective function (Section 3.2.1), the objective function may be defined as follows:

Minimize: Cost per mass unit of TCE destroyed (3.11)
Subject to: Specified constraints on decision parameters

In order to define a fitness function, we used the same cost formulation previously developed (Section 3.2.3). The mass of TCE destroyed in the system is a cumulative calculation of the mass destroyed over the simulation time period. The fate and transport

model performs this calculation. Therefore, the fitness function follows the same format as the objective function (3.11):

$$fitness = \frac{Cost_{actual}}{MassDestroyed} \tag{3.12}$$

where

$Cost_{actual}$ = The annualized cost of a specified system [$]

Mass Destroyed = Mass of TCE destroyed in simulation period [mg * 10]

The unit for mass destroyed was chosen as 10 mg by trial-and-error in order that the magnitude be comparable to magnitude of the cost.

With the new fitness function we used the model grid described in Section 3.3.1, with a 500 day simulation period, and results were tested using the model grid described in Section 3.5. The contamination source was set at a concentration of 5 mg/L in a 15-m wide area in the two aquifers. In addition, initial contaminant concentrations, also at 5 mg/L, were designated over the entire center region of the model grid (Figure 3.4) in both aquifers. These initial concentrations were used to ensure that no treatment well configuration had an advantage of being able to destroy more mass just based on proximity to the contaminant source. However, this advantage may still be seen at later times for certain configurations because concentrations near the source area will be higher. To help compensate for this, the duration of the simulation was set only to 500 days. The shorter duration also allows us to decrease computing time for the GA. Also, if the α values generated for the initial solutions were 180° or 0°, they were changed to 135° and 325° respectively (Appendix A).

In implementing the new fitness function (Equation 3.12), the same GA parameters as previously discussed were used (Section 3.4), and the same range of

values was used for the decision parameters, except that a diameter of 15 cm was used for the catalyst column.

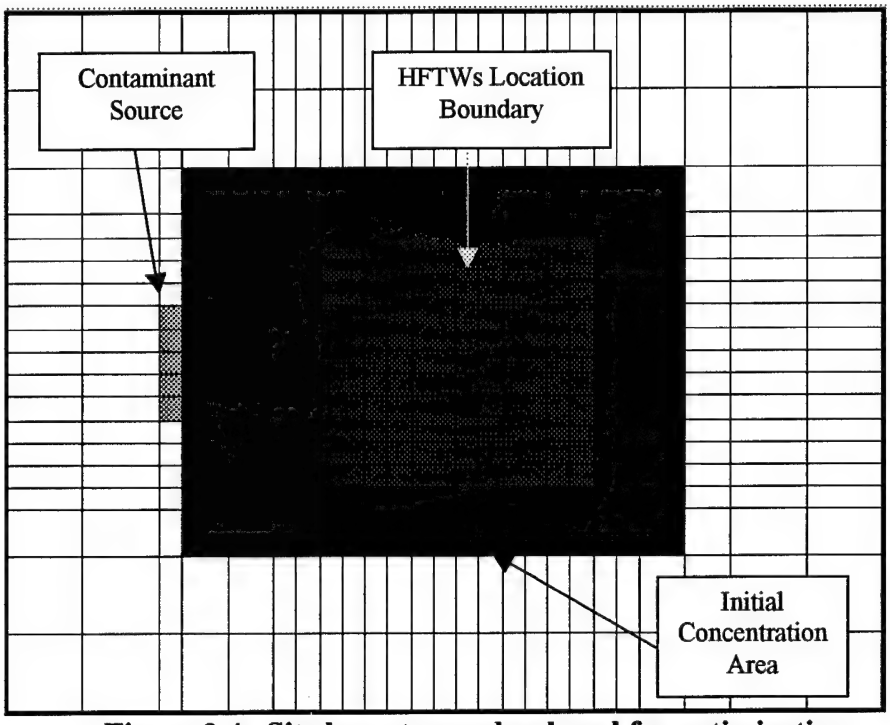

Figure 3.4. Site layout map developed for optimization.

4.0 RESULTS AND ANALYSIS

4.1 OVERVIEW

In this chapter, the results of the application of the genetic algorithm (GA) to the problem of optimizing palladium-catalyzed *in situ* destruction of TCE-contaminated groundwater are discussed. In the first section, a discussion of the results obtained using the cost per TCE mass destroyed fitness function are presented. In the second section, these results are analyzed.

4.2 RESULTS OF THE GA USING THE COST/MASS FITNESS FUNCTION

As previously discussed (Section 3.6), each simulation of the groundwater model was run for a duration of 500 days. Because an initial population was generated systematically (Section 3.4), the GA was expected to converge on a solution faster than if the initial population had been randomly generated. For this reason, and to save time, the GA was only run for 70 generations. Because of time limitations, only one full run of the GA was performed. This run took a total of approximately 86 CPU hours on a Pentium III, 650MHz processor, and 77 CPU hours on a Pentium III, 800 MHz processor.

The best solution was found by the GA on the 23rd generation and subsequent generations (up to 70) did not improve the solution. The "optimal" solution was:

$Q = 35.3\ m^3/d$

$d = 30.0\ m$

$\alpha = 360^\circ = 0^\circ$

$L = 2.33\ m$

The solution resulted in:

$Cost_{actual}$ = $3470

TCE Mass Destroyed = 39.0 grams

Fitness = ($Cost_{actual}$ / (TCE Mass * 100)) = 0.8898 ($ / 10 mgTCE)

where

$Cost_{actual}$ = The annualized cost of a specified system ($)

TCE Mass = Mass of TCE degraded (g)

100 = factor to convert mass from g to 10mg units

Figure 4.1 shows how the GA performed over the entire run. From the figure, we can tell that at least one of the individuals generated in the initial population had fitness value close to the "optimal" solution (1.081). The engineered parameters for the individual were: $Q = 81.6\ m^3/d$, $d = 28.3$ m, $\alpha = 326°$, and $L = 2.33$. As demonstrated in the discussion of cost (Section 3.2.3), one of the largest factors that affect cost is the length of the column, L, because palladium is very expensive and the energy required to pump water increases as L increases. Fortunately, some of our initial solutions, including the best solution of the generation, had the optimal L value. Because of the elitist technique employed (best solution is automatically passed to the next generation), most of the "best" solutions in subsequent generations had a value of 2.33 for L (Figure 4.2). Tests of the model using different well placement configurations, and varying L, show that the optimal value for L for all configurations tested is between 2 and 3 meters (Figure 4.3), which matches the "optimal" value obtained.

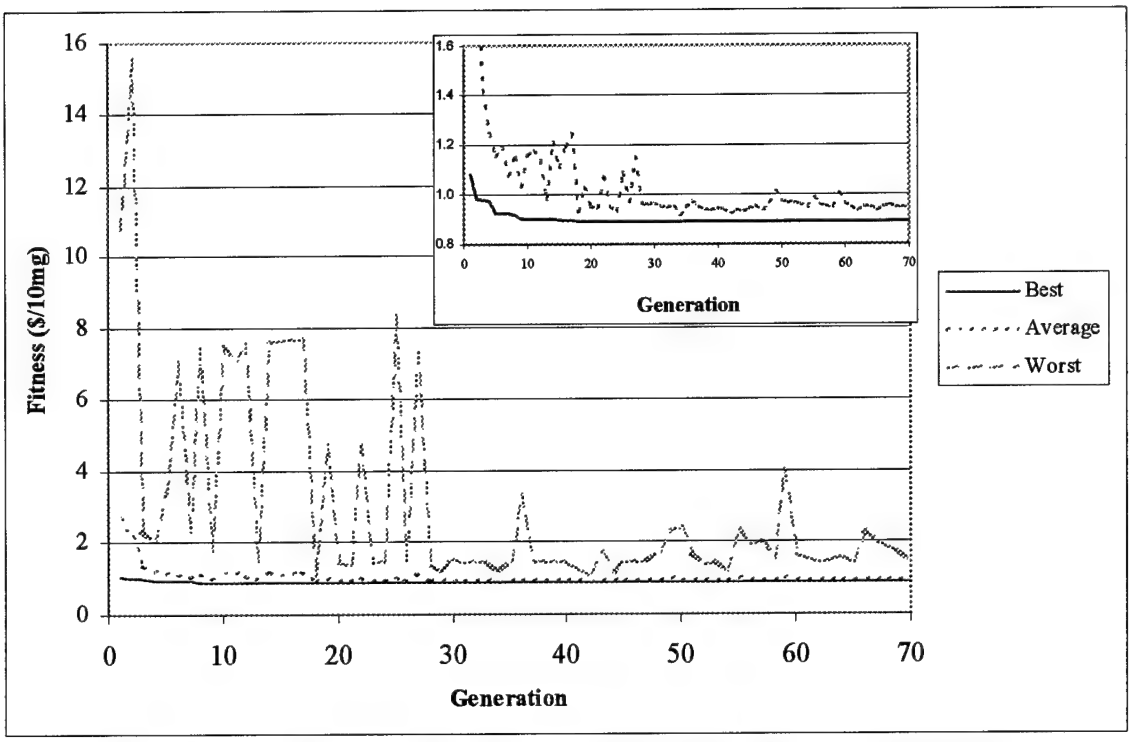

Figure 4.1. Best, average, and worst fitness function per generation.

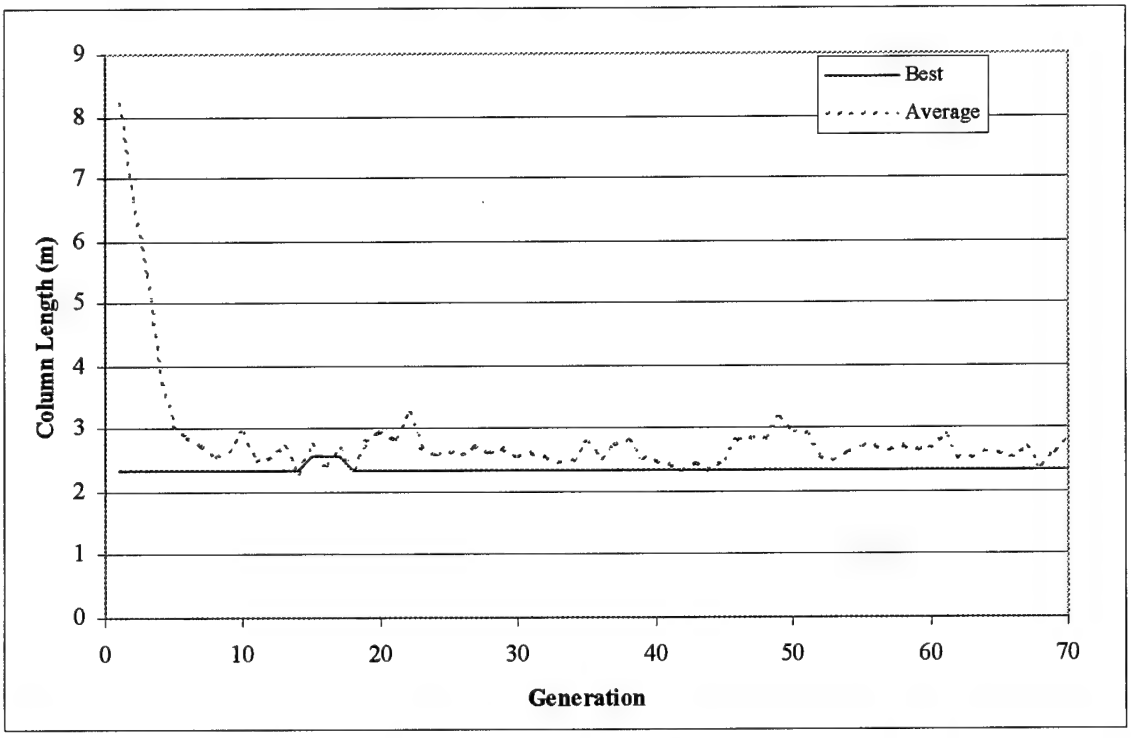

Figure 4.2. Column length of the best and average individuals in each generation.

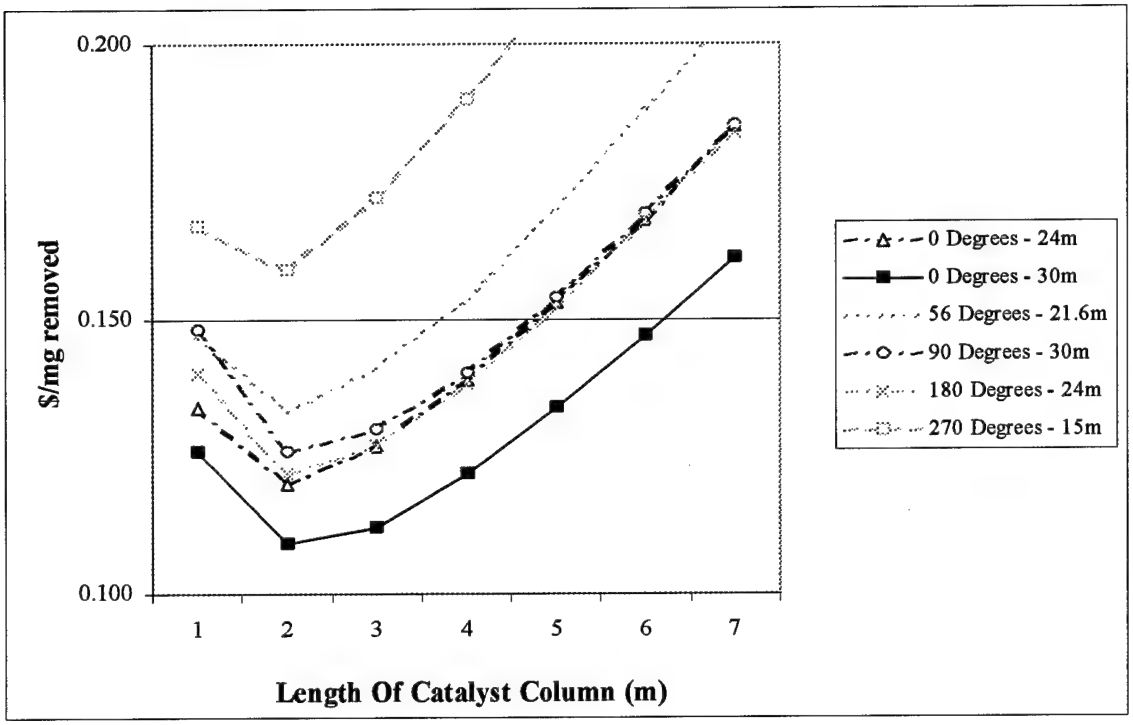

Figure 4.3. Cost per mass removed vs. length of catalytic column (365 day simulation).

Figure 4.1 also shows that the worst fitness value greatly varies over the duration of the GA run. The irregularity of this curve is due to mutations or crossovers that yield children with worse fitness than their parents. However, the frequency of these abnormally high fitness values decreases in the later generations because most of the solutions are nearing the optimal solution. Therefore, crossovers combine the genes of two good solutions, and if worse children are generated, these are not usually much worse. Because mutations are very infrequent, and not all mutations have negative effects on the individuals, mutations that generate significantly worse individuals are extremely rare. Therefore, the fitness of the "worst" individual also improves over the generations and converges to near the optimal value. If the GA did not allow for worse individuals to replace parents, we would always see the "worst" value converging towards the minimum fitness value, and its fitness would never deteriorate from one

generation to the next. The occasional bad fitness value for an individual in a population is also what causes the average value to slightly increase at times.

Using the higher resolution model grid to validate the optimal solutions obtained provided similar results to those obtained with the lower resolution grid. The total TCE mass removed was 39.4 grams, which is slightly more than the mass removal calculated using the original grid, and the fitness was 0.8787, which is slightly lower (better) than originally calculated. These improved results are to be expected because in the higher resolution grid there is less numerical dispersion of contamination, therefore slightly higher concentrations of contaminant are treated in the wells. Tests of other near optimal solutions generated by the GA showed that the "optimal" solution determined with the lower resolution grid was also superior when tested with the higher resolution grid.

4.3 ANALYSIS OF RESULTS

The results of the optimization show a configuration similar to that of the first optimization attempt, when we used the concentration-based fitness function. Using both the concentration-based and mass-based fitness functions, the value for L was similar (3.00 m and 2.33 m respectively). The reason why the two values are not exactly the same is that in the mass-based fitness function, there is not a requirement for a concentration downgradient. Violating the concentration requirement in the original fitness function penalized the fitness function, which in essence added "cost" for solutions that did not meet the concentration constraints. Thus, it was advantageous to increase the size of the column (even though capital costs increased) so that more mass would be degraded and the downgradient constraint could be met. A simulation of the

"optimal" solution from the concentration-based optimization yielded 32.56 g of TCE destroyed and a fitness value of 1.1159 $/10mg TCE, which is much worse than our "optimal" value obtained with the mass-based fitness.

Similarly, changing the value of Q affects both the cost and the performance of the system. Increasing Q decreases the residence time of the contaminated water in the catalytic column, but it increases the volume of water treated. An increase in this volume also means an increase in the mass of contaminant going through the column. Increasing Q also causes a minimal increase in the cost of the system (see Section 3.2.3). Likewise, decreasing Q has the opposite effects. In the case of the concentration-based fitness function, the optimal value for Q was much lower (15 m^3/d) than in the mass-based optimization (36 m^3/d). The lower value allowed for higher efficiency, so that the effluent concentration was adequate to meet downgradient concentration requirements. However, this low pumping rate did not allow for complete capture of the contaminant plume (Section 3.6). With the mass-based fitness function, the optimal pumping rate is determined as the point at which mass destruction is most cost effective.

Based on the system configuration, we can estimate the expected downgradient concentrations for both "optimal" solutions using an analytical approach (Equation 2.2). Such a comparison would give us an idea of the effect of changes in column length and pumping rate on overall contaminant destruction efficiency. These results can then be compared to results obtained numerically with the fate and transport model. If the inflow concentration into the upflow treatment well is 5 mg/L (a conservative estimate), we can determine the outflow concentration using the engineered parameters of our solutions and the volume of the reactor (Sections 3.2.3 and 3.6). For the parameters of the mass-based

solution, we would get a residence time of 42.3 seconds (4.899×10^{-4} d) for each column. Thus, the concentration leaving the upflow treatment well would be 2.07 mg/L (since the single-pass efficiency would be 58.6%). Assuming this concentration enters the downflow well located downgradient, its effluent would be 0.857 mg/L, or 857 μg/L. The overall efficiency of the system, comparing upgradient and downgradient concentrations in the aquifer being cleaned, would be 82.9%, (though it would be 0% for the other aquifer (Section 3.6)). A test of the system using the fate and transport model showed that the TCE concentration at an observation point in the lower aquifer downgradient of the downflow well was 0.148 mg/L (calculated without imposing an initial concentration in the treatment area and using a simulation duration of 1000 days). This concentration, which is well above the regulatory limits, would be highly penalized in the concentration-based optimization and would not be close to the optimal. However, the solution provides the most cost-effective configuration for mass removal. The reason for the difference between the analytically and the numerically obtained downgradient concentrations (0.857 mg/L and 0.148 mg/L, respectively) is that the analytical result assumes an inflow of 5 mg/L rather than 1.05 mg/L, which is the inflow concentration in the numerical model (determined by placing a monitoring well in the same area as the treatment well). The inflow concentration in the numerical model is just 1.05 mg/L for two main reasons. First, the well pumps water from all directions, and since there are no initial concentrations, the water just downgradient of the treatment well is clean. When downgradient clean water is pumped through the treatment well it dilutes contaminated water being pumped from upgradient. Second, dispersion of the contaminant as it flows

from the source area to the treatment well also contributes to the lower initial inflow concentration at the treatment well.

On the other hand, using the parameters obtained with the concentration-based fitness function, a residence time of 128.3 seconds (1.484×10^{-3} days) would result in each column. For this residence time, and assuming upgradient concentrations of 5 mg/L, we would expect an effluent of 345.6 μg/L and 23.9 μg/L in the first (downflow) and second (upflow) columns, respectively, corresponding to single-pass efficiencies of 93.1% and an overall (two-pass) efficiency of 99.5%. A test of this system using the fate and transport model resulted in a concentration of 8.2 μg/L at a well directly downgradient of the upflow well in the upper aquifer (using a simulation duration of 1000 days). As we can see, the solution obtained using the concentration based fitness function provides much higher overall efficiency. However, in both cases, one of the aquifers did not receive any treatment, and the contaminant plume was forced to travel around the re-circulation zone that was established between the two treatment wells (Section 3.6). In this aquifer, downgradient concentrations were an order of magnitude higher than downgradient concentrations in the aquifer where water was treated.

In any case, the combination of Q and L is what determines the single-pass efficiency of the system. As we can see in Figures 4.2 and 4.4, the average values for L and Q converged to their optimal values as the GA progressed through the generations.

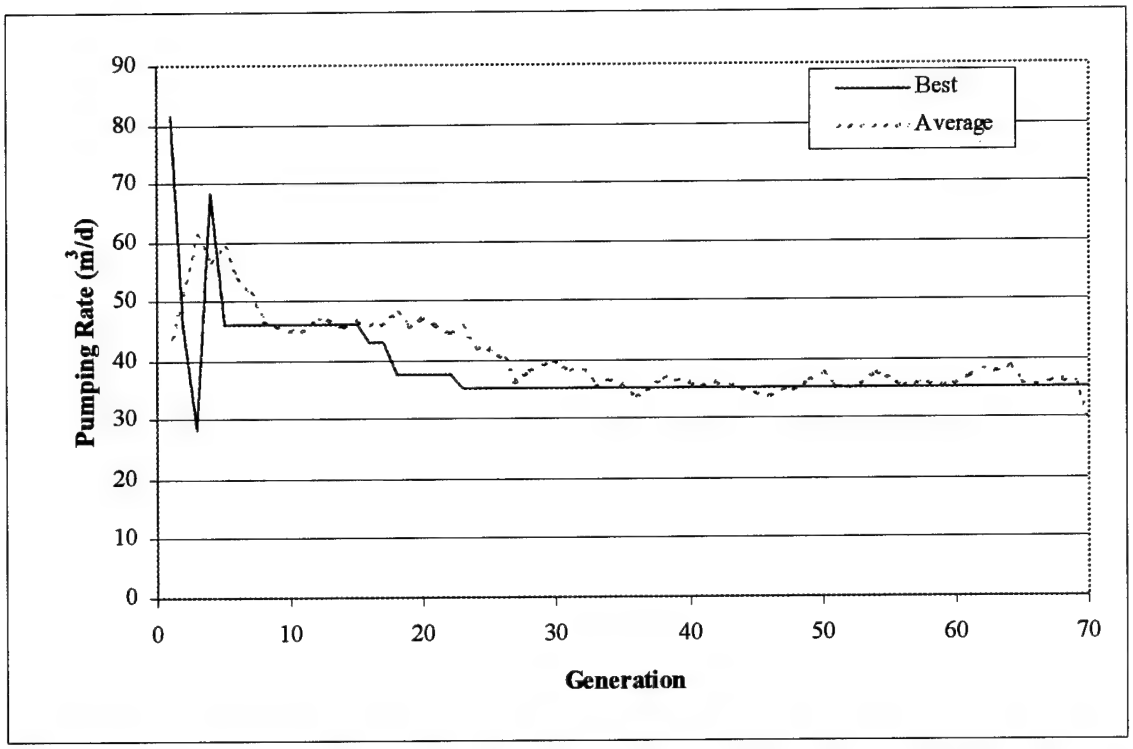

Figure 4.4. Pumping rate of the best and average individuals in each generation.

As previously discussed (Section 3.2.2), the distance between the wells, d, and the angle of regional groundwater flow relative to a line connecting the two pumping wells, α, determine the location of the treatment wells in our model grid. One of the main advantages of using an HFTW system is that the contaminated water is re-circulated through treatment zones. However, the amount of re-circulation, or interflow, between the wells is dependent on the location of the wells and their pumping rates. The placement of the wells in the optimal solution was 360°, or 0°. This means that the upflow well was placed closest to the source. Intuitively we might have thought the downflow well would be placed closest to the source (α = 180°) because the upper aquifer (Layer 2) is thicker than the lower aquifer (Layer 4) (8 and 6 m respectively), and therefore it contains more contaminant mass to be destroyed. However, whereas the lower aquifer is confined, the upper aquifer is overlaid by another layer (Layer 1), which

is partially saturated (Section 3.3.1). In our study, there was not a concentration source in this top layer, and the groundwater in it was clean. The downflow well pumps water from the upper aquifer, which causes drawdown near the well. This drawdown means that some water is pumped from the top layer into the downflow treatment well.

Model tests with α values of 0° and 180° were performed with and without a contaminant source in the top layer. Trials #1 and #2 show the aquifer conditions as tested in our study (no source in the top layer), and trials #3 and #4 were conducted with a contaminant source in the top layer. As expected, when there is a contaminant source in the top layer, the amount of mass degraded significantly increases for $\alpha = 180^\circ$ and remains fairly constant for $\alpha = 0^\circ$. This allows us to conclude that the downflow well causes enough of a drawdown in the upper layer to significantly reduce the average concentration entering the treatment well. Thus, for the mass-based fitness function, it is more desirable that the upflow treatment well be closest to the contaminant source ($\alpha = 0^\circ$).

Trial #	Layers with Cont. Source	Well Closest to Source (Angle)	TCE Mass Degraded (g)	Fitness ($/10 g)
1	2,4	Upflow (0°)	39.43	0.8787
2	2,4	Downflow (180°)	35.37	0.9796
3	1,2,4	Upflow (0°)	39.44	0.8787
4	1,2,4	Downflow (180°)	39.64	0.8742

Table 4.1. Fitness values and mass degradation comparisons for additional contaminant source area on water table layer (Q = 35.34 m3/d, L = 2.33 m, d = 30 m).

As mentioned in Section 3.6, an α value of 0° or 180° was expected because placing the pumping well as close as possible to the source area allows that particular well to treat water with very high concentrations, which provides more mass removal. To test this proposition, both wells were placed 18 meters away from the source, which is the distance to the source for the upflow well in the optimal solution, at an angle $\alpha = 90^\circ$,

using the same Q and L values from the optimal solution. The only parameter that was varied was the distance between the wells. As we can see in Table 4.2, the further apart the wells are placed, the higher the mass degraded. This indicates that the cost is maximized when interflow is minimized. The reason for this is because higher inflow concentrations into the catalyst column mean higher amounts of mass destroyed. As interflow increases, the amount of previously treated water being re-circulated through the system increases, thereby diluting the contaminant concentration in the well inflow. As a result, the optimal value for d obtained in the optimization was 30 m, which was the maximum value allowed in our study. In this configuration ($\alpha = 0^{\circ}$, d = 30 m), the amount of contaminant degraded is maximized by placing the upflow well close to the source and downflow well as far as possible from the upflow well. As discussed previously (Section 3.6), this configuration provides 100% interflow in the upper aquifer (from the upflow well to the downflow well), and minimal interflow in the lower aquifer (from the downflow well to the upflow well). Placing the wells at an angle of 0° minimizes interflow in the lower aquifer because the water must flow against the natural hydraulic gradient in order to be re-circulated. However, as previously demonstrated, such a configuration would result in violations of regulatory constraints downgradient of the treatment zone, especially in the aquifer that receives no treatment.

Distance (Meters)	Mass Degraded (grams)	Fitness ($/10 mg)
3	5.66	6.12246
6	12.77	2.71364
9	17.98	1.92686
12	21.59	1.60503
15	24.25	1.42885
18	26.15	1.32504
21	27.39	1.26497
24	28.38	1.22097
27	28.96	1.19668
30	29.48	1.17541

Table 4.2. Fitness values for well placement close to contaminant source at different well separation distances (Q = 35.34 m^3/d,$\alpha = 90^o$, L = 2.33 m).

Because this type of system seems to perform better with lower interflow between the wells, the ultimate solution would be to use multiple pairs of wells configured in a manner that both contaminated aquifers receive treatment. Different configurations could be used, such as a row of treatment wells (Figure 4.5) or two rows of treatment wells (Figure 4.6). Two rows of wells (Figure 4.6) essentially corresponds to alternating well pairs located using the optimal solutions obtained in this study ($\alpha = 0^o$ and $\alpha = 180^o$). These configurations of multiple pairs of wells would provide a larger effective capture zone in both aquifers and would allow interflow between the wells, so as to attain required overall efficiencies.

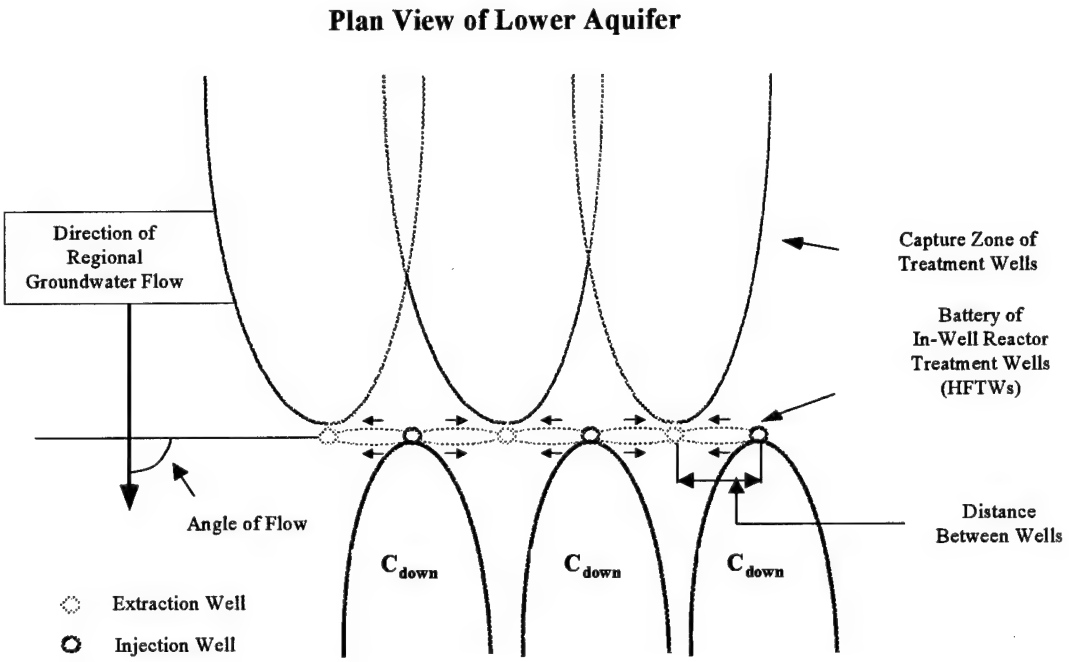

Figure 4.5 Row of multiple pairs of treatment wells

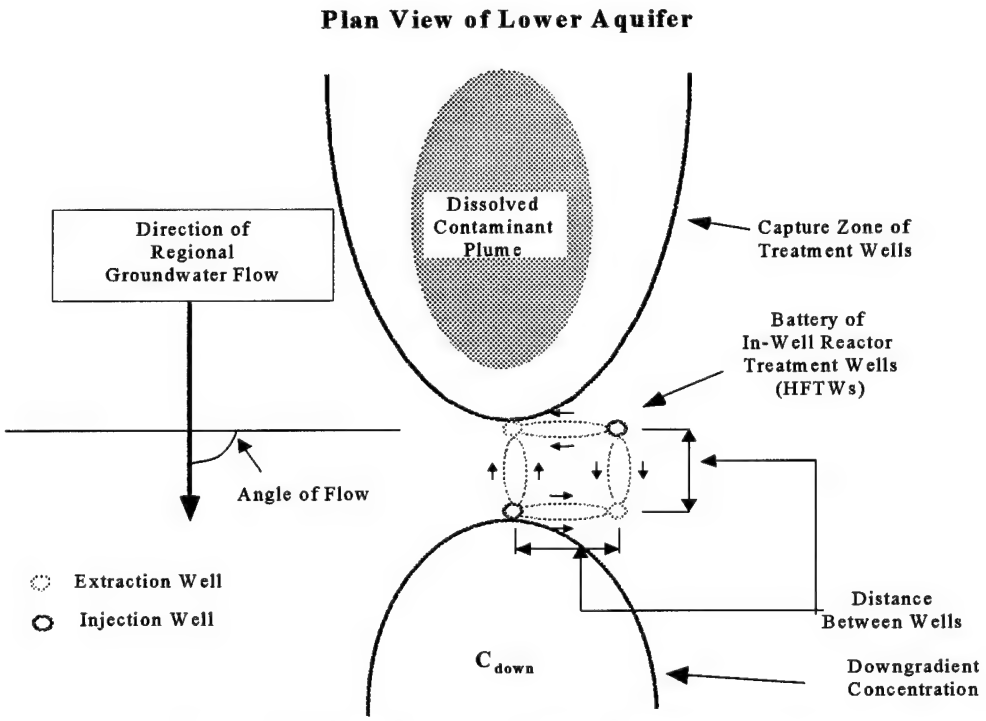

Figure 4.6. Two one-pair rows of treatment wells.

5.0 CONCLUSIONS AND RECOMMENDATIONS

5.1 SUMMARY

In this study, a method for determining how to best implement palladium-catalyzed *in situ* destruction of TCE-contaminated groundwater using a pair of horizontal flow treatment wells (HFTWs) was developed and applied to a hypothetical contaminated site based on conditions at Edwards AFB, CA. In the study, an objective function was developed and used to define a fitness function based on treatment technology cost and TCE concentration requirements downgradient of the treatment zone. A genetic algorithm (GA) was used to determine a technology solution that optimized the fitness function. The GA found an artificial solution for the problem that seemed to meet the constraints of the objective function, but in actuality did not provide adequate clean-up of the aquifer. Based on these results, new objective and fitness functions were developed in an effort to determine the most cost effective solution to remove contaminant mass from the aquifer. The solution arrived at using this approach, while resulting in minimized values of cost per mass removed, produced unacceptably high downgradient contaminant concentration levels.

The GA software that was used is a robust, user-friendly package that can be applied to a variety of objective/fitness functions and remediation sites with minimal user modifications of the software code. The contaminant fate and transport model used can also be modified to simulate different site conditions or treatment technologies. Together, these software packages can be used to enhance efforts to apply palladium-catalyzed *in situ* dechlorination using HFTWs to clean up chlorinated ethane-contaminated aquifers.

5.2 CONCLUSIONS

A GA will determine a solution that best meets the given constraints and fitness function. Using the two different fitness functions that were developed, we arrived at solutions that were similar. Note, however, that the solution that was optimal using one fitness function was very poor when evaluated using the other function. If the objective is to destroy mass as cost-effectively as possible, the mass-based function would provide a very good solution. However, if there are regulatory constraints, as in our study, a concentration-based function would provide a better solution.

GAs are effective optimization algorithms that can successfully be used to optimize groundwater treatment technologies. Even though the "optimal" solutions did not attain the desired downgradient concentration goals, the GA found solutions that optimized the fitness functions within the given constraints. Limiting the number of treatment wells to two did not allow for contaminant cleanup of both aquifers, which caused the bottom aquifer to receive no treatment. Also, the constraint placed on the simulation times produced spurious results, as the contaminant plume did not reach a steady state concentration at the downgradient monitoring wells within the time simulated. The simulation time was limited due to the excessive run time needed by the fate and transport model. The GA performed well in determining the best possible solution given these constraints as it ensured that the concentration downgradient of the treatment area was near the regulatory limits. Also, with the mass based fitness function, the GA demonstrated that it could find a very good solution that would minimize the cost per mass degraded.

5.3 RECOMMENDATIONS

Optimize an HFTW system using multiple pairs of wells. Placing multiple pairs of wells in a row over the entire length ($\alpha = 90^{o}$) of the contaminant plume would ensure the capture of the entire plume. However, the distance between the wells, pumping rates, and length of catalyst columns must be optimized in order to achieve the optimal interflow that provides a balance between cost effectiveness and achievement of downgradient regulatory concentrations constraints. However, other configurations, such as multiple rows, should also be explored. Perhaps, by placing wells in pre-determined locations and allowing the GA to select which wells will pump (and at what pumping rates), an optimal configuration could be determined. Based on the relative costs of buying palladium for the reactors and installing wells, it may be more cost effective to install a large number of wells with small catalytic columns, or a small number of wells with larger columns.

Incorporate catalyst deactivation and regeneration into the fate and transport model and optimize this model. Such an optimization would determine not only the optimal configuration of the wells, but also the most effective catalyst regeneration schedule. Catalyst deactivation is caused by natural compounds present in the groundwater and would vary depending on site geochemistry. Modeling this effect would enable the GA to determine optimal solutions for different sites and therefore ensure better system predictability upon implementation.

Apply the GA/fate and transport model to other geological sites. The site modeled in this study, although similar to Edwards AFB, is idealized. Optimization of

the technology at a variety of sites would further demonstrate the potential (and limitations) of palladium-catalyzed *in situ* dechlorination using HFTWs.

Develop a method to shorten the computational time requirements of the optimization. Because of the number of iterations involved in optimization using GAs, it takes a long period of time to generate one solution. One solution for this would be to use Artificial Neural Networks (ANNs) to take the place of the fate and transport model in the same manner as Rogers et al. (1995) (Section 2.7). Such an approach would involve much fewer runs of the fate and transport model, and would significantly decrease GA run time. It would also allow for higher resolution with the fate and transport model leading to more accurate results.

Parallelize the GA code. Running this study on a cluster of parallel computers as Garrett (1999) could have improved the overall runtime of the optimization. Garrett (1999) used twelve 200 and 250MHz computers, while we only used one computer at a time (divided between a 650 and 800 MHz computers). His first run of 100 generations on those computers took approximately 6 days, while ours took 6.5 days for 70 generations (equivalent of 9.7 days for 100 generations). However, Garrett's (1999) fate and transport model was more complex, requiring more time for each simulation. Had we used 12 parallel computers, we could have conceivably finished one run of the GA in less than two days.

APPENDIX A: INITIAL SOLUTIONS

The initial solutions generated to "seed" the GA were generated using a nested phase-center design. The purpose of this design was to ensure coverage of the entire solution space to prevent the GA from considering only one region of the solution space. A phase-center design consists of combinations of selected low, middle, and high values for each parameter. However, because of the large range of some of our parameters, we chose a nested design, and used two low values and two high values for each parameter. These were chosen approximately as 10%, 30%, 50%, 70%, and 90% of the maximum value for each parameter. A total of 50 initial solutions were generated. They consisted of the following (number of solutions in parentheses):

- all possible combinations of 10% and 90% values (16)
- all parameters set to 50%, except one at 10% (4)
- all parameters set to 50%, except one at 90% (4)
- all parameters set to 50% (1)
- all possible combinations of 30% and 70% values (16)
- all parameters set to 50%, except one at 30% (4)
- all parameters set to 50%, except one at 70% (4)
- all parameters set to 100% (1)

In the second GA run (using mass-based fitness function), if the initial solutions contained α values of 50% and 100% of the maximum (180° and 360°), these values were set to 135° and 325°, respectively. This was done to ensure the GA was not biased towards solutions containing these values from the beginning. Table A.1 lists the initial solutions generated.

#	Q	d	α	L	#	Q	d	α	L
1	8.34	2.25	33.75	1.31	26	25.03	6.75	101.25	3.94
2	79.96	2.25	33.75	1.31	27	62.58	6.75	101.25	3.94
3	8.34	21.75	33.75	1.31	28	25.03	16.50	101.25	3.94
4	79.96	21.75	33.75	1.31	29	62.58	16.50	101.25	3.94
5	8.34	2.25	323.44	1.31	30	25.03	6.75	253.13	3.94
6	79.96	2.25	323.44	1.31	31	62.58	6.75	253.13	3.94
7	8.34	21.75	323.44	1.31	32	25.03	16.50	253.13	3.94
8	79.96	21.75	323.44	1.31	33	62.58	16.50	253.13	3.94
9	8.34	2.25	33.75	12.69	34	25.03	6.75	101.25	9.84
10	79.96	2.25	33.75	12.69	35	62.58	6.75	101.25	9.84
11	8.34	21.75	33.75	12.69	36	25.03	16.50	101.25	9.84
12	79.96	21.75	33.75	12.69	37	62.58	16.50	101.25	9.84
13	8.34	2.25	323.44	12.69	38	25.03	6.75	253.13	9.84
14	79.96	2.25	323.44	12.69	39	62.58	6.75	253.13	9.84
15	8.34	21.75	323.44	12.69	40	25.03	16.50	253.13	9.84
16	79.96	21.75	323.44	12.69	41	62.58	16.50	253.13	9.84
17	79.96	12.00	180.00[a]	7.00	42	62.58	12.00	180.00[a]	7.00
18	8.34	12.00	180.00[a]	7.00	43	25.03	12.00	180.00[a]	7.00
19	44.50	21.75	180.00[a]	7.00	44	44.50	16.50	180.00[a]	7.00
20	44.50	2.25	180.00[a]	7.00	45	44.50	6.75	180.00[a]	7.00
21	44.50	12.00	323.44	7.00	46	44.50	12.00	253.13	7.00
22	44.50	12.00	33.75	7.00	47	44.50	12.00	101.25	7.00
23	44.50	12.00	180.00[a]	12.69	48	44.50	12.00	180.00[a]	9.84
24	44.50	12.00	180.00[a]	1.31	49	44.50	12.00	180.00[a]	3.94
25	44.50	12.00	180.00[a]	7.00	50	89.00	24.00	360.00[b]	14.00

Table A.1. Initial solutions generated for optimization effort.

[a] - denotes value changed to 135° in mass-based optimization
[b] - denotes value changed to 325° in mass-based optimization

BIBLIOGRAPHY

Ahlfeld, D. P., J. M. Mulvey, G. F. Pinder, and E. F. Wood. "Contaminated groundwater remediation design using simulation, optimization, and sensitivity theory, 1. Model development." *Water Resources Research, 24(3):* 431-441, 1988.

Aly, A. H. and R. C. Peralta. "Optimal design of aquifer cleanup systems under uncertainty using a neural network and a genetic algorithm." *Water Resources Research, 35(8):* 2523-2532, 1999.

Bäck, T. and D. B. Fogel. Glossary. *Evolutionary Computation 1: Basic Algorithms and Operators*. T. Bäck, D.B. Fogel, and T. Michalewicz. Bristol, United Kingdom: Institute of Physics Publishing, 2000.

Bäck, T., D. B. Fogel, D. Whitley, and P. J. Angeline. Mutation Operators. *Evolutionary Computation 1: Basic Algorithms and Operators*. T. Bäck, D.B. Fogel, and T. Michalewicz. Bristol, United Kingdom: Institute of Physics Publishing, 2000.

Bäck, T. Introduction to Evolutionary Algorithms. *Evolutionary Computation 1: Basic Algorithms and Operators*. T. Bäck, D.B. Fogel, and T. Michalewicz. Bristol, United Kingdom: Institute of Physics Publishing, 2000.

Bouwer, H. *Groundwater Hydrology*. New York, NY: Mc-Graw Hill, Inc., 1978.

Brown, T. L., H. E. LeMay, and B. E. Bursten. *Chemistry: The Central Science* (7th Edition). Upper Saddle River, NJ: Prentice Hall, 1997.

Carroll, D.L. D.L. Carroll's FORTRAN Genetic Algorithm Driver. Version 1.7. Computer Software ReadMe File. 1998.

Christ, J.A. A Modeling Study for the Implementation of *In Situ* Cometabolic Bioremediation of Trichloroethylene-Contaminated Groundwater. AFIT/GEE/ENV/97D-03. Wright Patterson AFB, OH, Air Force Institute of Technology, 1997.

Christ, J. A., M. N. Goltz, and J. Huang. "Development and application of an analytical model to aid design and implementation of *in situ* remediation technologies." *Journal of Contaminant Hydrology, 37*: 295-317, 1999.

Cieniawski, S. E., J. W. Eheart, and S. Ranjithan. "Using genetic algorithms to solve a multiobjective groundwater monitoring problem." *Water Resources Research, 31 (2)*: 399-409, 1995.

Culver, T. B. and C. A. Shoemaker. "Optimal Control for Groundwater Remediation by Differential Dynamic Programming With Quasi-Newton Approximations." *Water Resources Research, 29(4):* 823-831, 1993.

Dasgupta, D. and Z. Michalewicz. "Evolutionary Algorithms -- An Overview." *Evolutionary Algorithms in Engineering Applications.* D. Dasgupta and Z. Michalewicz. Berlin: Springer-Verlag, 1997.

Defense Environmental Restoration Program (DERP). *Fiscal Year 1997 Annual Report to Congress,* Washington, Government Printing Office, 1998.

Domenico, P. A. and F. W. Schwartz. *Physical and Chemical Hydrogeology* (2nd Edition). New York, NY: John Wiley & Sons, Inc., 1998.

Dougherty, D. E. and R. A. Marryott. "Optimal groundwater management, 1, Simmulated Annealing." *Water Resources Research, 27(10):* 2493-2508, 1991.

Eshelman, L. J. Genetic Algorithms. *Evolutionary Computation 1: Basic Algorithms and Operators.* T. Bäck, D.B. Fogel, and T. Michalewicz. Bristol, United Kingdom: Institute of Physics Publishing, 2000.

Feng, P. P. Modeling the Effect of Nonlinar and Rate-Limited Sorption on the Natural Attenuation of Chlorinated Ethenes. MS Thesis, AFIT/GEE/ENV/00M-04. Graduate School of Engineering and Management, Air Force Institute of Technology (AU), Wright Patterson AFB, OH, 2000.

Ferland, D. R. *In Situ* Treatment of Chlorinated Ethene-Contaminated Groundwater Using Horizontal Flow Treatment Wells. MS Thesis, AFIT/GEE/ENV/00M-05. Graduate School of Engineering and Management, Air Force Institute of Technology (AU), Wright Patterson Air Force Base, OH, 2000.

Fogel, D. B. Principles of evolutionary processes. *Evolutionary Computation 1: Basic Algorithms and Operators.* T. Bäck, D.B. Fogel, and T. Michalewicz. Bristol, United Kingdom: Institute of Physics Publishing, 2000.

Garrett, C. A. Optimization of *In Situ* Aerobic Cometabolic Bioremediation of Trichloroethylene-Contaminated Groundwater Using a Parallel Genetic Algorithm. AFIT/GEE/ENG/99M-02. Wright Patterson AFB, OH, Air Force Institute of Technology, 1999.

Goldberg, D. E. *Genetic Algorithms in Search, Optimization, and Learning.* Reading, MA: Addison-Wesley, 1989.

Guan, J. and M. M. Aral. "Optimal remediation with well locations and pumping rates selected as continuous decision variables." *Journal of Hydrology*: 20-42, 1999.

Harbaugh, A. W. and M. G. McDonald. User's documentation for MODFLOW-96, an update to the U.S. Geological Survey modular finite-difference ground-water flow model: U.S. Geological Survey Open-File Report 1996, 96-485, 56p.

Holland, J. H. *Adaptation in Natural and Artificial Systems.* Ann Arbor, MI: The University of Michigan Press, 1975.

Huang, J. and M. N. Goltz. "A Model of *In Situ* Bioremediation that Includes the Effect of Rate-Limited Sorption and Bioavailability." *Conference on Hazardous Waste Research Proceeding,* 1998.

LaGrega, M. D., P. L. Buckingham, and J. C. Evans. *Hazardous Waste Management.* New York, NY: McGraw-Hill, Inc, 1994.

Lowry, G. V. and M. Reinhard. "Hydrodehalogenation of 1- to 3-carbon halogenated organic compounds in water using a palladium catalyst and hydrogen gas." *Environmental Science and Technology, 33 (11):* 1905-1910, 1999.

Lowry, G. V. and M. Reinhard. "Pd-Catalyzed TCE Dechlorination in Groundwater: Solute Effects, Biological Control, and Oxydative Catalyst Regeneration." *Environmental Science and Technology, 34(15):* 3217-3223, 2000.

Mandalas, G.C. An Examination of Elements Necessary for Transfer of an Environmental Remediation Technology from Full-Scale Demonstration to Implementation. AFIT/GEEM/ENV/97D-17. Wright Patterson AFB, OH, Air Force Institute of Technology, 1997.

Mandalas, G.C., J.A. Christ, and M.N. Goltz. "Screening Software for an Innovative *In Situ* Bioremediation Technology." *Bioremediation Journal,* 2(1): 7-15, 1998.

Marryott, R. A., D. E. Dougherty, and R. L. Stollar. "Optimal groundwater management: 2. Application of simulated annealing to a field-scale contamination site." *Water Resources Research, 29(4):* 847-860, 1993.

Masters, G. M. *Introduction to Environmental Science and Engineering* (2nd Edition). Upper Saddle River, NJ: Prentice Hall, 1997.

Maymó-Gatell, X., T. Anguish, and S. H. Zinder. "Reductive Dechlorination of Chlorinated Ethenes and 1,2-Dichloroethane by 'Dehalococcoides ethenogenes' 195." *Applied and Environmental Microbiology, 65(7):* 3108-3113, 1999.

McCarty, P. L., M. N. Goltz, G. D. Hopkins, M. E. Dolan, J. P. Allan, B. T. Kawakamin, and T. J. Carrothers. "Full-Scale Evaluation of *In Situ* Cometabolic Degradation of Trichloroethylene in Groundwater through Toluene Injection." *Environmental Science and Technology, 32(1):* 88-100, 1998.

McKinney, D. C. and M. D. Lin. "Genetic algorithm solution of groundwater management models." *Water Resources Research, 30 (6):* 1897-1906, 1994.

McNab, W. W., R. Ruiz, and M. Reinhard. *"In-Situ* destruction of chlorinated hydrocarbons in groundwater using catalytic reductive dehalogenation in a reactive well: testing and operational experiences." *Environmental Science and Technology,* 34(1): 149-153, 2000.

Metcalf and Eddy, Inc. *Wastewater Engineering: Treatment, Disposal, and Reuse* (3rd edition). New York, NY: McGraw-Hill, Inc., 1991.

Merriam Webster's Collegiate Dictionary (10th Edition). Springfield, MA: Merriam-Webster, Inc., 1993.

Morgan, D. R, J. W. Eheart, and A. J. Valocchi. "Aquifer remediation design under uncertainty using a new chance constrained programming technique." *Water Resources Research, 29(3):* 551-561, 1993.

National Research Council. *Alternatives for Groundwater Cleanup.* Washington, D.C.: National Academy Press, 1994.

Reeves, C. R., and J. E. Beasley. Introduction. *Modern Heuristic Techniques for Combinatorial Problems.* C. R. Reeves. London: McGraw-Hill Book Company Europe, 1995.

Reeves, C. R. Genetic Algorithms. *Modern Heuristic Techniques for Combinatorial Problems.* C. R. Reeves. London: McGraw-Hill Book Company Europe, 1995.

Ritzel, B. J., J. W. Eheart, and S. Ranjithan. "Using genetic algorithms to solve multiple objective groundwater pollution containment problem." *Water Resources Research, 36 (5)*: 1589-1603, 1994.

Rogers, L. L. and F. U. Dowla. "Optimization of groundwater remediation using artificial neural networks with parallel solute transport modeling." *Water Resources Research, 30(2)*: 457-481, 1994.

Rogers, L. L., F. U. Dowla, and V. M. Johnson. "Optimal Field-Scale Groundwater Remediation Using Neural Networks and the Genetic Algorithm." *Environmental Science and Technology, 29 (5)*: 1145-1155, 1995.

United Stated Environmental Protection Agency (US EPA). *Use of Monitored Natural Attenuation at Superfund, RCRA Corrective Action, and Underground Storage Tank Sites,* Office of Solid Waste and Emergency Response Directive 9200.4-17P, 1999.

United Stated Environmental Protection Agency (US EPA). *New Final NPL Sites.* Excerpt from unpublished article. n. pag. http://www.epa.gov/superfund/sites/newfin.htm. 21 February 2001.

United States General Accounting Office (US GAO). *Environmental Clean Up: DOD's Implementation of the Relative Risk Site Evaluation Process*, Report GAO/NSIAD-99-25. Washington, Government Printing Office, 1998.

United States General Accounting Office (US GAO). *Superfund: Half the Sites Have All Cleanup Remedies in Place or Completed*, Report GAO/RCED-99-245. Washington, Government Printing Office, 1999.

Vasquez, J. A., H. R. Maier, B. J. Lence, B. A. Tolson, and R. O. Foschi. "Achieving Water Quality System Reliability Using Genetic Algorithms." *Journal of Environmental Engineering, 126(10):* 954-962, 2000.

Wiedemeier, T. H., M. A. Swanson, D. E. Moutoux, E. K. Gordon, J. T. Wilson, B. H. Wilson, D. H. Kampbell, P. E. Haas, R. N. Miller, J. E. Hansen, and F. H. Chapelle. *Technical Protocol for Evaluating Natural Attenuation of Chlorinated Solvents in Ground Water*, EPA Report 600/R-98/128, 1998.

Yoon, J. H. and C. A. Shoemaker. "Comparison of Optimization Methods for Ground-water Bioremediation." *Journal of Water Resources Planning and Management, 125(1):* 54-63, 1999.

VITA

Capt Manuel Fernández was born in Santurce, Puerto Rico. He was raised in Humacao, Puerto Rico until age 13, when he moved to Columbia, Tennessee. He graduated from Columbia Central High School in 1992, and entered undergraduate studies at the Georgia Institute of Technology in Atlanta, Georgia. He graduated with a Bachelor in Civil Engineering and was commissioned as an Air Force officer on 19 December 1996. In August 1999, he entered the Graduate Engineering and Environmental Management program of the School of Engineering, Air Force Institute of Technology.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From – To)*
20-03-2001	Master's Thesis	Aug 1999 – Mar 2001

4. TITLE AND SUBTITLE

OPTIMIZATION OF PALLADIUM-CATALYZED *IN SITU* DESTRUCTION OF TRICHLOROETHYLENE-CONTAMINATED GROUNDWATER USING A GENETIC ALGORITHM

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Fernández, Manuel, Captain, USAF

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 P Street, Building 640
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT/GEE/ENV/01M-02

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Center for Environmental Excellence/Technology Transfer Division
Attn: Maj Jeff S. Cornell
3207 North Road, Building 532
Brooks AFB TX 78235-5357 (210) 536-4329

10. SPONSOR/MONITOR'S ACRONYM(S)
AFCEE/ERT

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

13. SUPPLEMENTARY NOTES

14. ABSTRACT Conventional technologies for the treatment of groundwater contaminated with chlorinated solvents have limitations that have motivated development of innovative technologies. One such technology currently under development involves using palladium-on-alumina (Pd/Al) as a catalyst to promote dechlorination. Pd/Al catalyst may be used in-well as part of a re-circulating horizontal flow treatment well (HFTW) system. An HFTW system involves two or more dual-screened wells, with in-well reactors, to capture and treat contaminated groundwater without the need to pump the water to the surface.

In this study, objective and fitness functions, based on system costs and TCE concentration requirements, were developed to optimize a dual-well HFTW system with in-well Pd/Al reactors in a two-aquifer remediation scenario. A genetic algorithm (GA) was coupled with a three dimensional numerical model of contaminant fate and transport to determine optimized HFTW control parameters (well location, pumping rate, and reactor size). The GA obtained a solution within the specified constraints, but the solution was an artificial solution, as contaminated groundwater in one of the two aquifers received no treatment. Based on these results, new objective and fitness functions were developed in an effort to determine the most cost effective solution to remove contaminant mass from the aquifer. The solution arrived at using this approach, while resulting in minimized values of cost per contaminant mass destroyed, produced unacceptably high downgradient contaminant concentration levels. We conclude that by specifying that only two wells could be used in the HFTW system, we overconstrained the problem and that a multi-well HFTW solution is required.

15. SUBJECT TERMS

Contaminant fate and transport, horizontal flow treatment wells, catalyst, palladium, groundwater modeling, chlorinated ethenes, trichloroethylene, TCE, reductive dehalogenation, *in situ* treatment, aquifer remediation, optimization, genetic algorithm

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Professor Mark N. Goltz, ENV
U	U	U	UU	106	19b. TELEPHONE NUMBER *(Include area code)* (937) 255-3636, ext 4638

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39-18

Form Approved
OMB No. 074-0188